Methods, Models, and Conceptual Issues

An Invitation to Cognitive Science

Volume 4

edited by Don Scarborough and Saul Sternberg

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

Chapter 11

Skill Acquisition and Plans for Actions: Learning to Write with Your Other Hand

Patricia G. Lindemann and Charles E. Wright

Editors' Introduction

Suppose you are taking a Spanish course and you write the word *pato* (duck) in your notebook for the first time. If you think about it, it requires a lot of skill and coordination to move your fingers in the right direction and with the right amount of force. Even though you have never written this particular word before, it is not as though you are starting from scratch and learning to write all over again. You are writing letters that you have written many times as parts of other words. The point here is that a skill such as writing requires complex organization and planning of many smaller skills. Each word is composed of some of the letters one has learned to write, and each letter is composed of some of the pen strokes one has learned to produce.

How does your brain plan, organize, and coordinate the writing of an old or new word? You are probably not even aware of this complex coordination process, and most of it is difficult to observe in skilled performance because it occurs so swiftly. But if a person is put in a sufficiently novel situation, performance slows down, errors occur, and it becomes easier to identify and analyze the processes involved. This is just the strategy that Patricia Lindemann and Charles Wright pursue in their case study of what happens when right-handers are asked to write left-handed. The researchers used computer techniques to record and monitor the initially hesitant and clumsy strokes of the left-handed writing. Then, after the subjects gained extensive practice in writing particular words, Lindemann and Wright investigated what exactly they had learned. Did their newfound fluency apply to new words composed of practiced letters? To new letters composed of practiced strokes? To new letters composed of new strokes?

Their analysis reveals that at the level of the word and the letter, the left hand can make use of a lifetime of learning by the right hand, but that this is not the case at the level of the stroke. Their work also tells us that, despite our lack of awareness, handwriting skill is indeed organized in these levels of increasing specificity—that is, hierarchically. Lindemann and Wright consider exactly what a pen stroke is, what it might mean to perform a skill fluently, and how much practice it would take for the left hand to become as fluent as the right. Along the way they guide you in performing a small handwriting experiment of your own.

Chapter Contents

11.1 Exercising *Handwriting* to Introduce Our Topic

To start thinking about the topic of this chapter, take out a clean sheet of paper and do the following. First, in the upper left-hand corner of the paper write the word *handwriting* as you normally would. Now, below this write *handwriting* again, but this time use your nondominant hand, that is, the hand with which you normally do not write. Finally, tape or have someone hold the paper on the wall and write *handwriting* a third time, this time with your dominant hand but writing much larger than normal, so that the initial *h* is two to three inches tall.

Now take a look at the three instances that you wrote. Of course one is bigger than the other two, but otherwise do they generally look the same? Figure 11.1 shows an example of this exercise. The samples have been reproduced to be about equal in size so that it will be easier to compare their shapes. What this figure shows, and what most people who do

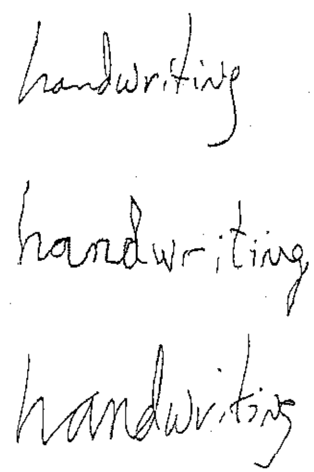

Figure 11.1
Samples of the word *handwriting* written with (from top to bottom) the right hand, left hand, and right arm. To facilitate shape comparisons, the reproductions were adjusted in size to have the same height for the ascenders. Originally the heights were 1.6 cm, 1.9 cm, and 4.9 cm, respectively.

this exercise generally find, is that, although there are clear differences—writing with the left hand is usually much shakier-looking for instance—there is a similarity in shape across these three productions. (For the remainder of this chapter we will be using "right hand" to refer to the dominant hand and "left hand" to refer to the nondominant hand.) This similarity goes beyond the obvious similarity that comes from the same word having been produced three times. In all three instances, the style of the writing tends to look similar (Wright 1990a, 1993). In fact, because of these shape similarities, handwriting experts are good at matching samples written by the same person using different body parts (effectors).

11.1.1 The Puzzle of Shape Similarity

At first, it might not seem remarkable that the same word produced by the same individual in three different ways would look so similar. To see why this similarity is important, note some of the factors that influence the control of writing in each case. One factor is the completely different physical makeup of each of the effectors. For example, different sets of muscles are used for writing by the right hand, right arm (used for the larger-sized writing) and left hand. A second factor is the amount of experience people have had writing with each effector: people write with their right hand every day, occasionally with their right arm (on blackboards and the like) and rarely, if ever, with their left hand. The shape similarity demonstrates that, despite substantial physical and experiential differences among effectors, there must be a common thread underlying these three ways of writing *handwriting*. We are interested in uncovering what this common thread might be.

Now think back for a minute and try to recall your sense of the process of writing each of the three samples. Were there differences in (a) how certain you were that you would succeed before you started, (b) how difficult it actually seemed while writing, or (c) how long the writing seemed to take? Although cognitive scientists are often appropriately wary about using introspective evidence such as this, it can sometimes provide hints about phenomena that can then be studied more objectively. In most cases, people doing this task find the process of writing with their right hand or their right arm subjectively quite similar. Having to write with the left hand is, however, often viewed with apprehension. Most people succeed well enough, but they usually have a definite feeling that writing with the left hand is harder, takes longer, and is somehow controlled differently.

Although people often feel that it is hard to write with their left hand because it is inherently less coordinated or weaker, this is rarely the primary problem. In fact, they use their left hand daily for a variety of skilled movements: for example, consider the important role of the left hand in touch-typing. In most cases, many of the problems initially encountered writing with the left hand can be overcome eventually through practice. To help you appreciate this, take another sheet of paper and write *handwriting* twenty additional times with your left hand.

11.1.2 Improving *Handwriting* with Practice

When you have finished your left-handed practice, compare your first and last attempts. Was your left-handed writing quicker after twenty repetitions? Can you see any difference in what you wrote? That is, do your last attempts look smoother, with more accurately formed letters?

Most people detect substantial improvements in their left-handed writing after even this minimal amount of practice. How much longer do you think you would have to continue practicing before your left-handed writing was as fluent as that of your right hand? What aspects of your left-handed writing would have to change in order for the writing to be considered fluent? Intuitively you may come up with answers to this last question that are similar to ours. We believe that fluent writing should be accurate and relatively fast. It should not be erratic and jerky like the initial attempts at left-handed writing tend to be. Instead, there should be a sense of steadiness, regularity, and control.

At the end of the chapter we will return to the question of how much practice is needed for the left hand to become as fluent as the right hand in writing. In the meantime, while you read the chapter, stop at each section heading and try writing *handwriting* five more times with your left hand. This will give you a chance to practice and improve your writing. By the end of the chapter you will be better able to assess what it would really take to learn to write fluently with your left hand.

(Here is a section heading, so practice writing *handwriting* five times here.)

11.1.3 The Goal of This Chapter

This chapter focuses both on the nature of the mental representations that underlie skills such as handwriting and on the way in which these skills are acquired. Why, for example, do the three instances of *handwriting* that you initially produced look so similar if, as introspection suggests, they were produced in different ways? Does the fact that they appear similar give us any clue about the underlying representation used to control the performance in each case? A specific question that this chapter addresses is, which parts of the information originally learned to make right-handed writing fluent can be incorporated into the process of learning to write with the left hand? Before grappling with these specific issues, however, it will help to consider the form of plans in general and something of the history of this concept in the area of motor control.

11.2 Plans and Planning

Consider planning and then carrying out a sequence of actions directed toward a specific goal. This might involve developing a high-level plan aimed at an ill-defined goal such as writing a term paper on the role of psychology in cognitive science, or it might be much more concrete, such as the plan you might construct for driving to work every day. In the second instance, the plan would be largely determined by the information you gain through daily experience with your commute. Your plan would probably include fairly detailed information about the roads you could take to work as well as the order in which to take them. In addition, your knowledge of where traffic is usually bad, which highways are under construction, and how much the tolls cost might all play a role in the development of your plan. As part of the planning process, you would explicitly make many decisions before executing the plan: what route you should ordinarily take, where you could stop to pick up a newspaper, and so forth. Other details you would leave undecided: how fast to go, whether or not to open the car windows, and so on. You might also include some contingencies in your planning: for instance, what to do if you discover that traffic is really bad on your usual route.

Once a plan exists, it has to be recorded in some way, both for immediate use and future reference. In a fairly simple or routine case (like the commute), you might just keep the plan in your mind. In a more complex case, such as planning for a term paper, you might want to set up a timetable with written goals, and perhaps use a calendar to ensure that the plan stayed on course.

(Remember, now is the time to practice your left-handed writing.)

11.2.1 Nature of Plans

In what ways can we characterize plans like those we have been discussing? What types of information should be included? At the very least, any plan must include some information about the actions to be executed and their order. The need for information about the actions is obvious. Without it, there would be no specification of what needs to be done. However, the ordering information is also critical because, in most situations, not just any ordering will do. In some cases, the required ordering is forced by the situation: for example, an object cannot be put down until it has been picked up. In others cases, the choice of order is dictated more by

considerations of efficiency than necessity: for example, you buy everything you need from one store at the same time, instead of making several trips. Although an ordering we produce in a given situation may not be the very best possible way to achieve a goal, given our constraints, we would expect it to be much closer to optimal than a random ordering. Thus one characteristic of planning is that it determines an efficient ordering of the steps necessary to achieve the goals.

A second characteristic of planning is that you can improve a plan as you gain relevant experience. This characteristic of planning adds to the efficiency with which a plan can be carried out. A substantial part of this improvement comes from being better able to take into account the dependencies between separate actions before the plan is executed. Consider the problem faced by a driver wishing to drive quickly on a poorly marked highway but having to exit at a particular exit. For a driver with experience making this particular trip, these requirements pose little problem. The driver stays to the left until about a half mile before the necessary exit and then starts moving to the right lane in anticipation of the upcoming exit. A driver less experienced with this route, however, may change to the slower lane prematurely or may have to change lanes abruptly when the exit is suddenly at hand. Similarly, if there is a choice between two different roads, the decision of a driver experienced with the route may be modified depending on the current weather conditions and how they will affect the roads.

Combining these characteristics, a plan can be seen as a selection and ordering of a set of actions intended to increase the efficiency with which the plan's goals can be achieved. Subsequent modifications of the plan can be made to help achieve goals even more efficiently. A successful plan leads to actions that usually seem to automatically anticipate the requirements imposed by the task and the environment.

11.3 Plans in Motor Behavior: Motor Programs

Within the domain of motor control, plans are typically referred to as "motor programs." Such programs, which take their name by explicit analogy to computer programs, are thought to be involved in activities as diverse as touching an object, hitting a ball with a bat, walking, running, pole vaulting, driving a car, writing with a pen, and producing speech. Thus motor programs are seen as the repositories for the accumulated information that underlies skilled, fluent activity, much as the plans discussed above (e.g. the commuting plan) are conceived of as repositories for the information needed to create an efficient, goal-directed action sequence.

Continuing this analogy, the general characteristics we hypothesized above for plans in general (that they contain action and sequencing information that can be modified through experience) could also be hypothesized for motor programs. This suggests a number of questions about motor programs, ranging from whether they actually demonstrate these general characteristics to how they could be structured and how the instructions they contain could be interpreted by the motor system. Clearly, a better understanding of the properties of motor programs for movement sequences, as well as of the planning and control processes that create and interpret them, could have far-reaching implications for learning and ultimate performance in many areas. In the specific case of handwriting, for example, better understanding could lead to the development of teaching materials that facilitate the process of learning to write and the development of better techniques for helping students to correct their handwriting problems.

11.3.1 A Movement-Specific Conception of Motor Programs

Although the term *motor program* has been used for more than twenty years to refer to the mental representations people use to perform skilled actions, its meaning has changed during this period. In the early usage of the term, a motor program was conceived of as a movement-specific, temporally structured collection of parameter values (i.e., force, velocity, duration, muscle sequencing) that, when transmitted directly to the motor system, would initiate and carry through a specific action. This conception was proposed largely to explain the development of coordination as skill in a complex movement increases.

Think about the process of developing your ability to perform a complex skill. When you first attempt a complex skill (for instance, writing with your left hand), you usually observe corrective processes operating: you make the first part (or submovement) of the overall movement, evaluate the result, make another submovement, reevaluate, and so on. The resulting performance is jerky, constructed of irregular submovements, and depends heavily on concurrent visual feedback (Pew 1966). There is a sense of knowing what you want to do but not really knowing how to carry out that intention. This might be something like having general directions for driving to a destination in an unfamiliar area, but not being sure of the details required to follow the directions: for example, what to look for as landmarks, how long to expect to travel between turns, or how to adjust to specific road conditions as you are driving.

As you continue writing with your left hand and your skill develops, the jerky submovements that you originally made become better coordinated with one another and appear to merge into larger, smooth strokes.

As this learning process proceeds, forming letters with the left hand becomes easier and requires less of your attention, allowing you to concentrate more on other aspects of the task. Once you have become an expert at writing *handwriting*, the motor program that results from your practice controls all the details of your newly acquired skill. But according to the original ideas about motor programs, the motor program that you have developed applies only to this single, specific task. It would not apply if, for example, you wanted to use your left hand to write *handwriting* larger or smaller than normal, or faster or slower than normal.

At this point, if you have remembered to do your left-handed writing practice, you will have written *handwriting* forty-six times with your left hand. In the process, you will have improved noticeably at this task. Now that you have developed some skill, see what happens when you temporarily interrupt this learning process for a quick experiment with your left-handed writing. Try, once or twice, writing *handwriting* with your left hand at one-half the speed you consider normal at this point. Now try writing it larger than normal. Was writing with either of these modifications as difficult as the first writing you did with your left hand at the start of this chapter? Most people find that these changes do not cause great difficulty or much deterioration in their performance. For a contrast, also try writing *skyscraper* with your left hand. Although *skyscraper* has one fewer letter than *handwriting* and the two words share several letters, most people find it much harder to write *skyscraper* for the first time than to write *handwriting* smaller/larger or faster/slower than usual.

(Continue writing *handwriting* normally at each section heading. Please do not practice other words until we ask you to again at the end of the chapter.)

11.3.2 Generalized Motor Programs

Observations similar to those just considered above prompted theorists to propose a more elaborate role for motor programs. In this view, a motor program specifies a class of similar movements rather than just a single skilled movement. As we observed above, such a class might include movements that vary in size or speed. A generalized motor program, because it controls a class of related movements, must undergo a transformation before being converted into muscle commands: variables contained in the generalized instructions, which allow delayed specification of at least some movement parameters, are replaced by specific values appropriate for the particular movement to be made. In the handwriting example, if writing speed were represented as a motor program parameter that had not yet been assigned a value, a value for writing speed would be specified before the movement could begin.

On an intuitive level, one strong appeal of the generalized motor-program hypothesis is that it explains how we produce novel variants of previously learned movements. We simply change the value of variables in the program. In contrast, according to the older concept of specific motor programs, we must use trial and error to create a new program for each new element in a class of similar movements.

Logic and perhaps your experience writing *skyscraper* suggest that there must be limits to the generalization that is possible. Take the commuting example considered earlier. If your car was being repaired and you had to borrow a friend's car, probably little about your commute would change, although you would be making somewhat different movements inside the car to control it. By contrast, it seems unlikely that a commuting plan that worked adequately in New York City would be of much help guiding you to work the first day of your new job in Los Angeles. This contrast suggests that route topology is an integral part of the commuting plan that is not generalized, although the plan is sufficiently abstract that it does not include all details of car control. (But would things be different if your friend drove a car with a standard transmission and you knew only how to drive a car with an automatic transmission?)

Thus the hypothesis of the generalized motor program creates an empirical and a theoretical challenge. Empirically, the challenge is to determine what happens when we learn a new skill that is a modification of an existing skill. Do we simply assign values to the parameters in a generalized motor program or do we generate a new motor program through a process of trial and error? Theoretically, the challenge of the generalized motor-program hypothesis is to create models of the computations necessary to generalize particular movement parameters that are consistent with other known regularities of motor performance.

11.3.3 Effector-Independent, Generalized Motor Programs

Later in this chapter, we will focus on the specific claim, made by some as part of the generalized motor-program hypothesis, that motor programs can be *effector-independent*. In the broadest interpretation of this claim, an effector-independent, generalized motor program, once established for a particular movement, can be used to control performance of that movement by any of a number of muscle-joint systems (*effectors*) in the body, although perhaps with some loss of precision. This claim evolves from the idea that specification of particular muscles and joints is not a necessary part of a motor program, but rather, as with other aspects of generalized motor programs such as amplitude and timing, the choice of effector may be represented by variables that are replaced by specific values only at the time of the movement. This is like saying, in the commuting example, that the route plan does not depend on whether you are driving your car or

that of a friend. Some people, including the authors (see Wright 1990a), have argued that, because of the very different geometries of the various muscle-joint systems in the body, this claim is implausible, at least in its most radical form.

Compare, for example, the joints and activating muscles used for handwriting—primarily the fingers and the wrist—with those used to write in a vertical orientation, as on a chalkboard—primarily the elbow and the shoulder. The hand, including the fingers, is an intricately organized web of bones and muscles. The fingers can be independently oriented at the knuckle, and they can bend quite smoothly because of their jointed construction. The hand itself can rotate up and down and side to side at the wrist and can even be formed into somewhat different shapes. By contrast, the elbow-shoulder combination has many fewer ways in which it can move (*degrees of freedom*). Although neither the elbow nor the shoulder has a simple structure, for writing on a chalkboard, they act essentially as hinges. These differences in joint organization and function are so large that it is hard to see how detailed, low-level plans for movements made by one of these joint combinations could be easily transformed so that they could be used to control the same movement made by the other combination. And yet, the finger-wrist combination is used for normal handwriting while the elbow-shoulder combination is primarily responsible for larger writing such as on a chalkboard, two related activities often suggested as candidates for control by a single, effector-independent motor program.

11.3.4 Evidence for Effector Independence in Handwriting: Shape Similarity

The evidence cited for the claim of effector independence in motor programs for handwriting consists largely of anecdotal demonstrations. Although Lashley (1942) was apparently the first to present and interpret this sort of demonstration, two later studies are usually cited. Merton (1972) compared writing produced by the fingers and wrist with writing ten times larger produced by the arm. Raibert (1977) displayed a sample in which his subject wrote the palindrome "Able was I ere I saw Elba" with the right hand, with the right arm (with the wrist splinted to eliminate movement), with the left hand, with the pen held in the teeth, and with the pen strapped to the right leg. These demonstrations are similar to the writing samples you produced at the outset of the chapter; they show a dramatic similarity in the shape of handwriting produced with different effector systems. Many see this similarity as evidence that a common motor program was used to control the writing in each of these examples.

11.3.4.1 A Problem Interpreting Shape Similarity

Consider what subjects take to be the task when instructed to write with their left hand, or what you take to be the task when you practice writing *handwriting*. Most likely, the goal is to produce writing that is as similar as possible to that produced under normal conditions. Thus the aspect of the performance the experimenter examines for evidence of invariance is the same aspect as the subjects use to judge their performance. Keep in mind that motor control is accomplished by a complex, flexible system. That such a system is capable of finding a solution for the task of producing a similar writing trajectory with a different effector tells us, we believe, more about the flexibility of this system than about the structure of motor programs.

11.3.4.2 Other Types of Evidence

This insight about interpreting shape similarity suggests that, when using handwriting as the object of these investigations, it might make more sense to look for evidence about effector independence in aspects of the performance other than the trajectory shape. In an earlier study, author Wright (1990a) took this approach. That study confirmed the overall shape similarity of writing produced with the right hand, right arm, and the left hand, as in figure 11.1, using more systematic, quantitative techniques than had been used in earlier work. More important, Wright 1990a also demonstrated that the similarly written shapes were produced in very different ways by the two hands. For example, writing with the left hand was substantially slower and exhibited a different decomposition of letters into strokes, as well as increased variability, and small but systematic changes of shape from that produced by the right hand. These results probably coincide with your experience writing *handwriting*. In addition, these results quantitatively demonstrate our previous observation that left-handed writing is less fluent than right-handed writing by demonstrating that it is slower and more variable.

The observation of these differences led to the conclusion that, to the extent that writing with the right and left hand share a common motor program, this representation must be quite abstract, incorporating little more than overall shape information and almost no detail about how to actually produce the writing. Thus one possible conclusion from Wright 1990a is that the right and left hands are not controlled in handwriting by a single, effector-independent motor program unless the motor program is very much more abstract than the generalized motor-program hypothesis would suggest. As we will see in what follows, however, this conclusion may be based on an overly simplistic conception of the way that information is represented in a motor program.

11.4 Hierarchical Representation of Plans

Although not considered in our earlier discussion of plans, it is possible that planning information exists at a number of different levels. Once again, think about the example of commuting to work. In that situation, your overall goal of traveling from your home to your place of work might constitute the highest level of planning. At the next level down, you might have more detailed plans for each part of your route (e.g. a "driving on Main Street" plan, a "driving on the interstate" plan, a "parking the car" plan). At an even lower level, there might be plans for performing the different actions required by each of these middle-level plans. For example, the "driving on the Interstate" plan could include lower level plans for "passing another car" and "paying a toll." Finally, there might be plans for controlling speed in cars with and without an automatic transmission. In this way, plans can be thought of as having a hierarchical structure with details specified at the lower levels and more global control specified at the higher levels (Greene 1972; Rosenbaum 1985; Saltzman 1979).

Clearly, a potential advantage of the hierarchical organization is that some of the more specific, lower-level plans may be generic. For example, the driving on Main Street" subplan could also be useful for organizing trips to the grocery store. Or the plan for "paying a toll" might apply equally well on the interstate or at a toll bridge on a secondary road.

11.4.1 From Reaching to Playing Shortstop: Development of Hierarchical Plans

As an example of the development of a hypothetical hierarchical motor representation, consider how an infant learns to reach through space and to grasp an object when it touches its hand. What might this learning consist of? How are these newly acquired movements represented so that they can be repeated? An infant's separate attempts at reaching and grasping can be thought of as two separate movement components at the lowest level in a hierarchy. Later, when the infant refines and combines these activities so that reaching for and grasping an object becomes a single, fluid motion, this would indicate the emergence of skilled control at the next level up in a control hierarchy. Interestingly, the process of combining skilled reaching and grasping submovements involves subtle changes in each of the components. At the same time, it appears that the skill for the component submovements is not lost when they are combined in a larger skill because they can still be produced separately. Later, at a still higher level of representation, the reaching and grasping combination may be further integrated with locomotion to allow a child to grab objects that are out of reach. Finally at the highest level in the hierarchy, these same activities may be further refined and combined with other separately learned sequences to allow an infielder in baseball to automatically (because there is no time to create a new plan) and fluidly dive for a hard-hit ground ball, catch it, roll, and come up throwing the ball to first base.

This type of analysis raises some interesting questions. How does being integrated into a higher-level activity change the components? For example, when an infielder makes a diving catch, how is the catch component different from when it is incorporated into catching a line drive or a fly ball? In what sense are complex, learned activities made up of previously learned components? What might components at different levels consist of? One focus of this chapter will be exploring these types of questions. Using handwriting as an example, we will look at what might be represented at different levels.

When thinking about the hierarchical representation of a motor program, it is important to keep in mind that we are referring here to the form of the mental representations and not to the neurophysiology of the motor system. This distinction is important, because the neurophysiological system for motor control also possesses a hierarchical structure (briefly outlined in an appendix to this chapter). Although it is quite plausible that some of the levels of the representational hierarchy that we have posited can be mapped onto the neurophysiological hierarchy, this mapping is probably not straightforward and, at the present time, is not well understood.

11.4.2 Hierarchy in the Motor Programs for Handwriting

The experiment that we will describe is based on the hypothesis that a motor program for handwriting has a hierarchical structure. As with the hierarchical plans we have just discussed, we conceive of this one as having several levels ranging from the highest (perhaps the intention to write a word or phrase) to the lowest (perhaps the specific muscle and joint commands to the wrist and fingers) with probably several levels in between. We suggest that at the highest level, information is effector-independent: the same intention can serve any effector. At the lowest level, information is effector-specific: the details of how to move the right hand must be different from the details about how to move the left hand. We will assume that somewhere in the middle levels of the hierarchy there should be a division such that information above that level is general enough to use an effector-independent representation, while information below that level is specialized enough to require an effector-specific representation.

For the purposes of this experiment, we will consider just three of the possible levels in the hierarchy controlling handwriting. The highest of these levels controls the production of particular, complete words. The second controls the production of the individual letters that make up each word. The third and lowest level controls the production of the particular strokes that make up each letter. Of course there may be additional levels, such as a higher level of representation for phrases or entire sentences, or a lower level in which explicit commands are given to particular joints and muscles.

11.4.3 Strokes in Handwriting and the Analysis of Tangential Velocity

Although what we mean by the word and letter levels probably seems clear, what we mean by the stroke level may be less obvious. Many researchers who study handwriting identify strokes as the basic unit in motor programs for handwriting (e.g., Edelman and Flash 1987; Hollerbach 1981; Viviani and Cenzato 1985). As we will see, the *kinematics* of written trajectories (the term refers to motions independent of the forces that produced those motions) appear strongly segmented. The points of segmentation separate pieces of the written trajectory that correspond closely to what people identify when asked to mark the strokes in their writing and correspond to what children are taught when they first learn to write. Finally, Margolin and Wing (1983) have described a brain-damaged male patient with a syndrome known as "apraxic agraphia",[1] whose major difficulty appears to be producing some, but not all, handwritten strokes from memory, even though he can produce a full range of nonwriting movements, which suggests that the handwriting strokes may be represented separately in the brain.

According to the conception that strokes are basic units of handwriting production, each letter is made up of a concatenated sequence of strokes. In an idealized version of this concept there is no overlap between strokes, and thus there are identifiable points where one stroke ends and the next stroke begins. In reality, because there may be some overlap at the transition from one stroke to the next, the identification of a specific boundary point between strokes can only be, at best, an approximation.

One way to see the segmentation of writing into strokes is to examine pen tip speed during the writing process. Although we often think of the pen, in skilled handwriting, as moving smoothly on the paper with a relatively steady speed, this is more of an illusion than reality. Panel A of figure 11.2 shows the x, y trajectory for the letter q written with the right hand. Panel B shows the *tangential velocity* (i.e., the writing speed) at each moment in time as the letter in panel A was produced. Notice how the tangential velocity increases and decreases over time contrary to the intuition that writing speed is steady.

We use local minima in tangential velocity to identify potential stroke boundaries. In single movements of a pen, from a specified starting position

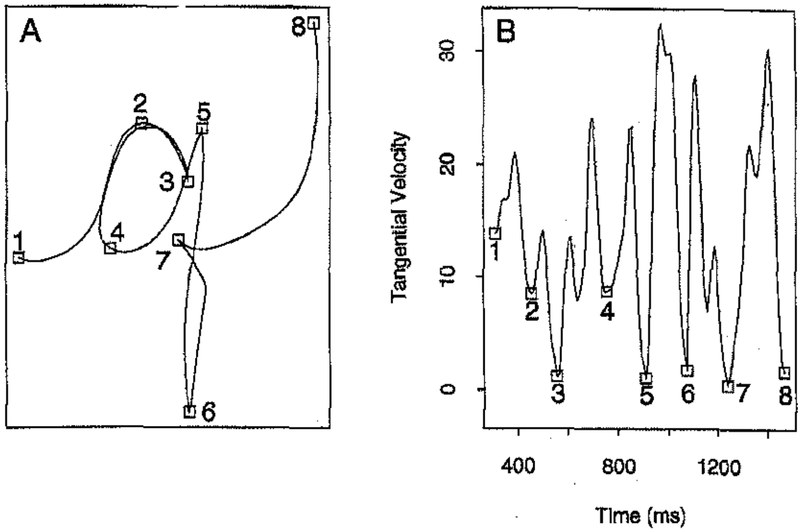

Figure 11.2
Illustration of the decomposition of the letter *q* into strokes based on local minima in the tangential velocity profile. Panel B, on the right, shows how the tangential velocity of the pen tip varies over time; panel A shows the *x*, *y* trajectory followed by the pen. The numbered squares in both panels mark analogous points.

to a target position, pen speed starts slowly, peaks in the middle of the movement, and then decreases again as the movement is completed. Thus, in handwriting, when strokes are produced in succession, we expect to see a slowdown in writing speed at the junctures where one stroke ends and the next begins. These slowdowns appear in the tangential velocity profile as local minima; these local minima can therefore be used to identify stroke boundary points.

There are eleven minima in the velocity profile shown in panel B of figure 11.2. Of these, only the eight marked with small boxes are taken as stroke endpoints delimiting the seven strokes in *q*. We do not interpret the remaining three velocity minima as stroke endpoints because they do not occur reliably across repetitions of *q* by this writer. The eight numbered boxes in the two panels of figure 11.2 label the same stroke boundaries in both panels.

It is possible that some of the strokes, identified in this way, are not separate units in the subject's internal representation of this letter. In general, tangential velocity slows at points of high curvature (Viviani and Terzuolo 1980), and some tangential velocity minima might simply reflect large curvature within a single unit. Despite this ambiguity, for our present purposes we have chosen to treat all of the strokes identified in this way as equivalent units.

11.4.4 Hierarchical Structure of Motor Programs: Why Does it Matter?

The assumption of hierarchical structure introduces several complications into the study of motor programs. If a motor program were simply a single, nonhierarchical set of instructions, then it would be natural to ask whether a motor program exists for a particular movement task and, if so, whether it is generalized in specific ways (e.g., across sizes or effectors). With the assumption that motor programs have a hierarchical structure, these questions no longer make as much sense. First, take the question of whether or not a motor program exists for a particular movement. It is difficult to know how to answer this question because, as skill develops, learning presumably is distributed across the levels in a hierarchy. Furthermore, skill may develop unevenly, with details completely specified for some levels of a task (where expertise has been achieved) but not at others (where learning is still taking place). When a child is first learning to write, for example, situations will arise in which the child knows how to produce several strokes but does not know how they can be combined to produce a letter, just as when you first wrote with your left hand, there were undoubtedly times when you had knowledge of the strokes necessary to produce a letter, but did not know how to produce them left-handed. Given a hierarchical structure, it seems more appropriate to ask about the representations at each level than about merely the existence of a motor program. Similarly, when we ask about generalization, it is important to recognize that the information at separate levels of the hierarchy may not have exactly the same type of representation. Information may be generalizable at some levels but not at others. Therefore, it probably does not make sense to ask whether a motor program is generalized. Instead, we should ask which of its levels are generalized.

What happens, then, as skill develops? In this conception, there is a continuum of change. As parts of the motor program are practiced, they become increasingly automatized and require less "attention" to produce; the resulting movements are more stereotyped and efficient. When these movements are produced in an appropriate context, they mesh smoothly together within that context and appear fluent. At the other extreme are parts of the motor program that have not been automatized (or, if automatized, that are inappropriate for the current task environment). In this case, the subject must constantly be aware of and make decisions about how to make a movement or continue an action sequence; skill acquisition is an extended, gradual process of automating the control structures in a plan so that they can handle increasingly large chunks of the task more of the time. Again, in view of this sort of gradual, distributed development of control autonomy and efficiency, it probably does not make much sense to ask whether a task is controlled by a generalized motor program. Instead, we need to ask what proportion of the learning ultimately possible at each level has been achieved at a given point in practice.

11.5 Studying Motor Programs: Two Approaches

11.5.1 Planning

Although our conception of motor programs as hierarchical suggests that the questions asked about motor programs may need to be reformulated, it is still useful to see how motor programs have been studied in the past. One approach that has been used to verify the existence of motor programs is to look for evidence of planning. Because motor programs involve planning and this planning is done before the planned movements begin, it is reasonable to expect knowledge of the entire movement sequence to influence the selection and production of each of its individual actions. For example, when getting out of the car, if you realize that you need to unlock the door of your house, then you can suppress the response of putting your keys away so that they will be out and ready when you get to the door. Sternberg and his colleagues have argued that one type of evidence for planning is that the production of actions throughout a skilled sequence depends on the number of actions in the sequence. Their data (summarized in Wright 1990b) show that, for both rapid speech and typing, the more actions there are in the sequence, the more time it takes to produce each action. If this observation applied only to the last action in a sequence, it could reflect a less interesting process such as fatigue, but this slowing applies to all the actions in the sequence, suggesting that the process of producing each action is affected by information about the entire sequence.

11.5.2 Transfer of Learning

The planning evidence does not provide much insight about the internal structure of motor programs. In the rest of this chapter we will use a different approach, measuring to what degree learning to produce one set of actions *transfers* to production of a second set of actions, in order to focus on the hierarchical structure of motor programs for handwriting. One objective is to determine at which level effector-specific information begins to be included in the control hierarchy (motor program) for writing and which levels of the hierarchy contain specifications that are generalized across effectors. Specifically, we will attempt to determine at which levels information must be relearned when learning to write with the other hand and at which levels there is effector independence (or in this case *hand* independence).

Posing the question in this way may seem to contradict the conclusion we reached earlier that writing with the right and left hands must not be controlled by a hand-independent motor program. Our consideration of the implications of hierarchical structure suggests, however, that another interpretation of those results is possible. Consider that the subjects in Wright's experiment (1990a) wrote with their left hand for only about one hour. Perhaps experience writing with the right hand led them to develop motor programs that contain hand-independent information at some levels, but because of limitations at lower levels specific to the left hand, the left-handed writing by these subjects could not benefit from these hand-independent representations.

To help clarify this idea, think back to the commuting example. Imagine someone who has developed a commuting plan over time, but knows only how to drive a car with an automatic transmission. What would happen if this individual were suddenly faced with the prospect of driving a stick-shift car to work? As observers, we might conclude that the previous commuting plan was "car-specific" because initially, at least until our hapless friend has mastered the intricacies of a clutch, all our friend's attempts at driving will be largely ineffective. If we were to watch again a day or two later, however, we might observe our friend driving to work as efficiently as ever and, if anything, with a little extra zip. Clearly, our initial hypothesis about the commuting program being car-specific was somehow mistaken.

Returning to handwriting, suppose that a motor program contains substantial hand-independent, middle-level information for the control of handwriting, but that at a lower level, additional hand-specific information is required. Because left-handed writing is rarely practiced, the low-level learning, specific to the left hand, is necessary before the more abstract, hand-independent information at the higher levels of the control hierarchy can usefully guide left-handed writing. As a result, initial writing with the left hand appears (as in Wright 1990a) to share only the most abstract specification of the trajectory shape with writing by the right hand. With additional left-handed practice it is possible that more detailed, hand-independent elements of the motor program developed for writing with the right hand could guide the left.

Some support for this position comes from a case study reported by Castiello and Stelmach (1993). They studied a left-handed male patient several years after his left hand was severed in an accident. After the accident the patient had learned to write with his (nondominant) right hand. Some years later, he obtained a prosthesis for his left hand and also became proficient writing with it. Once he had learned to write well with the left-hand prosthesis, this patient had acquired the unusual ability to write skillfully with both his right and (prosthetic) left hand. Castiello and Stelmach obtained samples of this subject's nondominant, right-handed writing and his prosthetic left-handed writing. Using the measures suggested by Wright (1990a), they compared these samples and reported similarities not only in the trajectory shapes but also in how those shapes were produced. These results are consistent with the hypothesis that fluent writing is controlled in similar ways across hands and suggest that fluent writing by the two hands shares a hand-independent representation at a low level in the hierarchy of the motor program.

11.6 Learning to Write with the Left Hand

The experiment reported here re-creates, under more controlled conditions, your experience as you first learned to write *handwriting* with your left hand and then tried to write *skyscraper*. We had two goals as we ran this experiment. The first was to learn more about the levels of the hierarchy we have hypothesized for the motor representation that underlies handwriting. The second was to determine at which levels in this hierarchy new learning is required for our subjects to be able to write fluently with their nondominant hand and, conversely, at which levels of the hierarchy the representation is hand-independent.

How could we investigate the representations that our subjects would develop for left-handed writing? The first step was for our subjects to start learning to write left-handed. To this end, they practiced an initial set of words until they had gained some left-handed writing ability. After our subjects had practiced this initial set of words, the critical step was to introduce them to new words and to see how fluently they were able to write these new words relative to the initial practice set. As we will describe in more detail shortly, if writing fluency on the new and old word sets was similar, this would suggest that aspects of the representation were hand-independent.

In our experiment, we carefully structured the words that our subjects learned to help us evaluate performance at the word, letter, and stroke levels. As you have been practicing your left-handed writing, you have been participating in a limited version of our experiment, and the (very limited) words you have been practicing have been structured accordingly. When we had you switch from *handwriting* to *skyscraper*, that change required you to do something new at all three levels of the hierarchical control representation. You had to make adjustments at all three levels because we were asking you to write a word that you had not

written previously with your left hand that included both letters and strokes that you had also not written previously with your left hand. We did this to maximize the difference between the word you had been practicing and the new word that you had not practiced.

To meet our goals for this experiment, we wanted to be able to look separately at transfer of learning for these three levels. This would allow us to isolate the effects at each level. How could we structure the words we introduced to make this possible? One way would be to introduce changes at each of the levels sequentially. First, we could introduce new words that differed from the initial practice set only at the word level—that is, the words would be new, but the letters used and strokes used would remain constant. This would effectively isolate the word-level representation. After determining whether the representation at the word level is hand-specific or hand-independent, the letter and stroke levels could be examined in the same way. In the following sections, we will see how this works, as we examine the learning process for writing in more detail.

11.6.1 Learning to Write with the Left Hand: The Number 1 *Team*

Think back to the point before you began writing with your left hand and imagine that you have to learn to write the word *team* fluently. According to our analysis, this involves acquiring information at least at the word, letter, and stroke levels. This is illustrated schematically in figure 11.3. To begin with, you must know the identity and serial order of the letters that make up the word: here *t*, followed by *e*, followed by *a*, and ending with *m*. This information constitutes what we have labeled the "word level" in figure 11.3. In addition to this sequencing information, the word level may contain the details that control production of the linking strokes that connect the letters in a word. These links differ, depending on the shape of the letters being connected (e.g., an *e*-to-*a* link differs from an *o*-to-*a* link because the letter *e* ends at the writing line while the letter *o* ends above the writing line). The word-level information is, however, hardly sufficient. In addition, you must know how to produce each of the four letters with the left hand. This is the information contained in the *letter* level.

Logically, the representations at the letter level could consist of complete specifications of the motions used to produce each letter or, even more specifically, the series of muscle activations required to produce those motions. As we discussed in section 11.4.4, however, there are good reasons to believe that the specifications at the letter level are in terms of strokes, which then constitute yet a lower level of organization.

If we take the stroke level as given, then, just as the word level consisted largely of a specification and ordering of the letters to be produced,

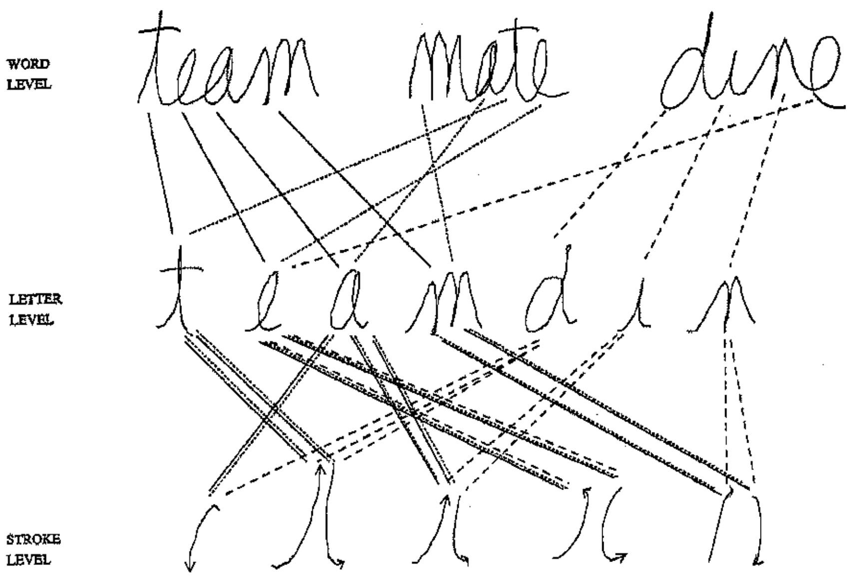

Figure 11.3
Hypothetical, hierarchical decomposition of the words *team, mate,* and *dine* into component letters and strokes. The lines connecting the word to the letter level and the letter to the stroke level are drawn in three styles (solid, dotted, and dashed) corresponding to each of the three words. The presence of a line is meant to indicate how, within the context of a particular word, elements that make up the unit at one level in the hierarchy are produced by generic units at lower levels in the hierarchy.

the representations at the letter level would consist primarily of a specification of the strokes and their sequence within a letter. And just as the word level might contain information concerning the links between the letters that make up a word, the representation at the letter level might include detailed specifications that facilitate the joining together of the set of strokes required to produce a particular letter.

Finally, according to this analysis, there is the *stroke* level of the representation. This level contains the specifications used to produce each of the stroke movements hypothesized to be the lowest level units of production in this representation.

We have tried to make it clear in this description that each of these levels might involve hand-dependent information that has been specialized so that the resulting movements can be produced fluently only by a particular hand. At the same time, it is conceivable that the representation at each of these levels might be hand-independent. Both possibilities appear plausible, even at the stroke level. For example, a representation at the

stroke level that was hand-dependent might consist of specifications of which muscles to activate, their amplitudes, and timings. (In keeping with our discussion in section 11.3.2, the specifications of amplitude and timing could be generalized so that the overall size and duration of the writing could be changed just before it was to be produced, but the specification of which muscles to activate would be a permanent part of the representation.) Alternatively, the representation developed at the stroke level could be hand-independent. As Keele (1981) has noted, Hollerbach's coupled-oscillator model of handwriting (1981) is an example of a mechanism in which the stroke-level representation is hand-independent.[2]

No matter how the levels of this hypothesized control hierarchy are organized into hand-dependent and hand-independent representations it seems that this organization must include some hand-dependent representations. This is because the left-handed writing and right-handed writing are controlled differently (Wright 1990a; described in section 11.3.4.2). If the writing is controlled differently, at some level the representations that are guiding the writing process must be specific to the particular writing hand. Assuming that there are some hand-dependent representations, how then might the process of learning to write left-handed be accomplished? According to our analysis, developing fluency in left-handed writing would likely involve creating or improving the relevant left-hand-specific representations.

As with learning any complex skill, learning to write a first word fluently with the left hand will take time, and performance will improve with practice. Panel A of figure 11.4 shows how a measure of performance might improve as a subject practices writing a word with the left hand. One measure of performance that could be used here is writing time. Later, we will discuss other measures that may be as useful. In writing time, as in golf, improvement results in a lower score.

11.6.2 Learning to Write with the Left Hand: Two Moves to *Mate*

Consider now that, having learned to write the word *team* fluently, you are required to learn a second word, *mate*. How will the performance for *mate* compare with that for *team*? Panel B of figure 11.4 shows four possible learning curves for performance after the switch to *mate*. Compare these to the original learning curve for practicing *team*, shown in panel A. Which (if any) of the continuations in panel B is most plausible depends on which levels of the hierarchical representation are hand-independent and also depends on the nature of the learning process.

In some ways, writing *mate* is to writing *team* as writing *skyscraper* is to writing *handwriting*. In both situations, we are comparing a word that has been practiced with the left hand (*team* or *handwriting*) to a word that has

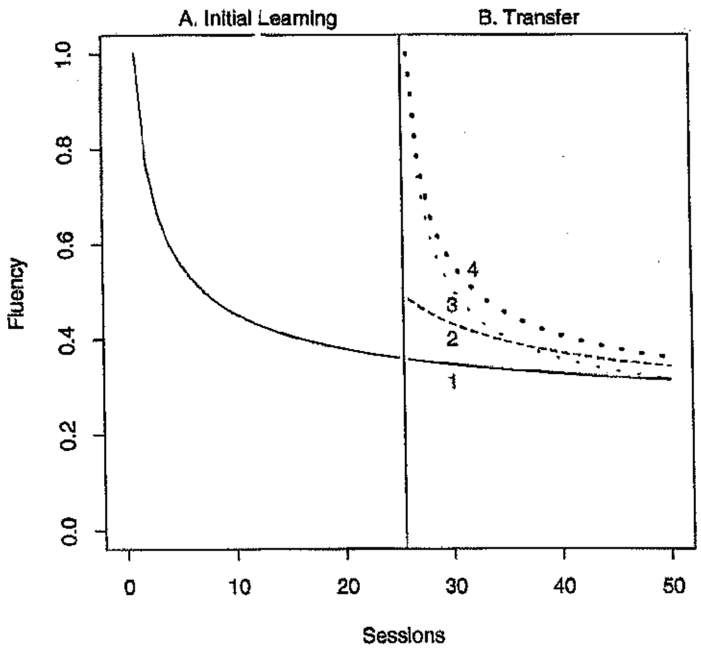

Figure 11.4
Plots of a power function for hypothetical data based on learning to write two words with the left hand. The solid curve in panel A displays the course of initial learning. The four curves in panel B illustrate possible outcomes when the word being practiced is changed. The solid line (1) illustrates the case in which there is no degradation of performance because the transfer is complete. The heavy dotted line (4) illustrates the case in which there is no transfer and the subsequent learning process is identical to the initial learning. The light dotted line (3) illustrates the case in which there is no transfer but the subsequent learning process benefits from the original learning and is faster. The dashed line (2) illustrates the case in which there is partial transfer of the original learning.

not (*mate* or *skyscraper*). One possibly important difference between these situations is that, as shown in figure 11.3, *mate* is composed from the same four letters as *team*, just arranged in a different sequence. Whether this difference is important depends on the truth of two things. First, because you have already practiced the necessary strokes and letters, you now know how to make those strokes and letters with your left hand; therefore it is possible that at the letter and stroke levels you can now function as efficiently for *mate* as for *team*, even the first time that *mate* is written. Second, because of your prior experience writing with the right hand, it is possible that you have a hand-independent representation for *mate* at the word level. If both of these suppositions are correct, then we would

Table 11.1
Possible outcomes for writing performance with the change from practicing *team* to *mate* (words constructed by rearranging the same letters)

		Letter and stroke level learning for *team* are word-independent: Yes	Letter and stroke level learning for *team* are word-independent: No
A hand-independent representation for *mate* exists at the word level based on prior experience with the right hand	Yes	A No change in the progression of writing time over practice	B Not relevant
	No	C An intermediate result that is hard to predict	D Complete relearning needed: writing time starts at its prepractice level

expect no deviation in the learning curve for writing speed with the change from learning to write *team* to learning to write *mate*. This outcome is illustrated by the curve 1 in figure 11.4B.

This outcome depends on choices between a pair of alternative hypotheses that are related to the suppositions above, as shown in table 11.1. Note that only the second of these pairs of alternatives, the one that distinguished between the rows in table 11.1, involves hand independence. The first pair of alternatives concerns the word specificity of learning at the letter and stroke levels—that is having learned to make a particular stroke or produce a particular letter in the context of one word, can it then be produced the same way in other words? The outcome we considered in the previous paragraph corresponds to the cell labeled "A" in table 11.1.

The other outcome for which there is a clear-cut prediction is that in the lower right cell, outcome D. The prediction for this outcome is curve 4 in figure 11.4B. This curve illustrates a complete failure of the learning produced, while practicing *team*, to generalize to the task of writing *mate*. Because of this failure of generalization, learning for *mate* must be started as if there had been no prior experience writing with the left hand. We would expect the data to fall on this curve 4 only if, as table 11.1 suggests, both of two conditions were met. First, at the stroke and letter levels, the prior learning with the left hand is completely word-specific. Second, at the word level, there is no hand-independent representation for *mate*.

Curve 3 in figure 11.4B represents a variant of the outcome just described that is often suggested. Here again, the representations are such that we are dealing with the lower right cell in table 11.1. Thus the learning process must be restarted at the beginning when the word being written is switched. In this variant, however, the learning is faster because of the prior experience learning to write *team*; although there is no direct generalization of learning from one word to the next, there is an improvement in the learning process.

We can ignore outcome B in table 11.1, because, although it is logically possible it makes no sense given what we know about handwriting.[3] This leaves outcome C, the lower left cell in table 11.1. Curve 2, the final learning function in figure 11.4B, represents this case. In outcome C, prior experience writing *mate* with the right hand is not generalized through a hand-independent representation to the left hand, but the skills learned at the letter and stroke levels writing *team* with the left hand are generalized. In this case, the performance for *mate* is initially somewhat better than that for *team* written for the first time, but worse than that for *team* written just prior to the switch. Just how much better or worse depends on issues about which we have no insight.[4]

11.6.3 Learning to Write with the Left Hand: Three to *Dine*

Let us assume that the observed learning curve for the switch from *team* to *mate* was curve 1 in figure 11.4B and is consistent with outcome A, described in the previous section and table 11.1. Based on this outcome, we infer that representations at the word level are hand-independent. By assuming this is so, we can concentrate our analysis of a second switch of materials on the letter and stroke levels. Thus, having mastered *team* and *mate*, suppose that your next task is to learn to write the word *dine*. This word presents new challenges similar to some of those you encountered when you switched from writing *handwriting* to writing *skyscraper* because *dine* is not simply a rearrangement of the letters in the previous two words. In fact, only the letter *e* is common to this and the previous two words. However, as figure 11.3 illustrates, it is possible that the strokes required to produce *dine* are the same as those in the words *team* and *mate*, although of course in *dine* they are arranged differently. (You might be skeptical about our claim that the strokes are the same across all three words. We will discuss how we justify this claim in the next section.)

Our analysis of the possible outcomes for this switch is summarized in table 11.2, which is similar to table 11.1. Notice, however, that because we are assuming that there are hand-independent representations at the word level, the word level is not mentioned in table 11.2. Instead, the two alternatives that distinguish the rows involve a claim of hand independence at the letter level. Similarly, the two alternatives for the columns involve the claim that there is a stroke level of representation and that prior learning at the stroke level is not word-specific.[5] Once again, the two most interesting results are those associated with outcomes A or D because these are the ones that we can interpret with the least ambiguity.

Table 11.2
Possible outcomes for writing performance with the change from practicing *mate* to *dine* (words with the same strokes but different letters)

		Stroke-level learning for *team* and *mate* is context-independent	
		Yes	No
A hand-independent representation for *d*, *i*, and *n* exists at the letter level bases on prior experience with the right hand	Yes	A No change in the progression of writing time over practice	B Not relevant
	No	C An intermediate result which is hard to predict	D Relearning of strokes and letters is needed: writing time starts at its prepractice level

Outcome A is again associated with a learning curve that does not change at the point of the switch between words. This outcome should occur only if there are hand-independent representations at the letter level for *d*, *i*, and *n*, and if the stroke-level learning, based on experience writing *team* and *mate*, is not letter-specific. If both conditions are met, then it should be possible to write *dine* as fluently on the first attempt as *mate* on the last.

If outcome D is correct—that is, the letter level is not hand-independent and prior learning at the stroke level does not generalize across letters—then, as you first start to write *dine*, there is no advantage at the letter or stroke levels of having previously practiced *team* or *mate*. This suggests that the learning curve would be like curve 4 in figure 11.4B and that performance would be as bad as it was originally for *team*. This result may seem surprising, given that there was no change in performance with the transition from *team* to *mate*. But consider that the hand-independent, word-level representation for *dine* should not be any more available or helpful than those for *team* and *mate* were before you had any left-handed writing practice. And, unlike the transition from *team* to *mate*, according to outcome D none of the prior learning at the letter and/or stroke levels is available as you first start writing *dine*.

Outcome C again leads to intermediate results. In this case, the necessary learning at the stroke level has occurred during the prior practice with the left hand, but there are no hand-independent representations at the letter level. Because of this failure of generalization at the letter level, we expect that the performance curve for outcome C will be between those of outcomes A and D.

To summarize this section on learning to write with the left hand, we started with an assumed three-level, hierarchical representation of motor programs for handwriting. If this assumption is correct, then the logic just outlined shows how, by manipulating the sequence of materials for subjects learning to write with their left hand, we may be able to identify properties of the representation at each of these levels by looking at particular patterns of the learning across practice. At each level of the hierarchy—word, letter, stroke—representations could be either hand-specific or hand-independent. Furthermore, the learning at each level could be context-specific or context-general. As subjects progress from the initial practice materials, to new words that contain no new letters and then to new words that do contain new letters, the resulting pattern of learning will help us to understand how general the representation is at each level.

11.7 Preparations: Identifying Generic Strokes, Characterizing Learning Curves, Methods and Design

Before we present the experiment based on the logic just described, we need to describe several techniques we used and the experimental design and methods.

11.7.1 Identifying Generic Strokes: G-Strokes

At various points in the discussion above, we suggested that letters and/or strokes can be generalized across words. For letters, what we mean probably seems clear. We all know, for example, what the letter *l* is and looks like. Thus it seems reasonable that having learned to write an *l* by practicing one word, you could use this experience when you need to write an *l* in some other word. A moment's reflection demonstrates that this may not actually be as straightforward as it seems. This generalization cannot simply be tied to the strokes in a particular letter; for example a capital *L* at the beginning of a sentence is written very differently from the small *l* we started thinking about. A similar, but more subtle problem, is that the *l* in *ol* starts out quite differently in most people's cursive writing from the *l* in *el*. With some refinements, however, the underlying idea is still comprehensible and, we hope, still plausible.

Thinking about strokes in an analogous fashion involves, however, an additional complication. Superficially, it is not obvious that writers produce the same strokes in the contexts of different letters. There is no standard labeling scheme for strokes that gives us the sense that there are strokes that are, at least nominally, identical across letters. The problem here is similar to that encountered by those who study speech. For speech, phonologists have over time identified a limited set of *phonemes*, sounds that are nominally identical across words. Similar analyses have also been done for handwriting but, so far as we can tell, less formally. Thus for example, letters are often introduced, to children learning to write, in "families" related by the perceived similarity of their component strokes.

What we are about to describe is a procedure we have developed to identify, in different letters, strokes that are similar both in their shape and in the way that they are produced. For want of a better name, we call these "generic strokes" or "g-strokes" for short. We need a reasonable procedure to identify g-strokes for two reasons. First, in the design and analysis of our experiment we use g-strokes to balance materials across conditions (this will be discussed later). Second, and more important our analysis suggests that a limited set of appropriately identified g-strokes are the generalized building blocks that are combined to produce all of the letters. For example, consider again the cursive letter *l*. It is made up of one upstroke and one downstroke. Now think about writing the letter *b*. In most people's handwriting, the letter *b* begins with two very similar strokes. According to our analysis, the same two g-strokes are used in both letters.

The process of identifying the g-strokes began with the right-handed writing for each subject. We analyzed each subject's writing individually because we were by no means certain that the decomposition of letters into g-strokes would be the same for all subjects. Our use of right-handed writing for this analysis was somewhat risky because it could result in the identification of g-strokes that were inappropriate for left-handed writing. However, because there was no way to base this analysis on each subject's left-handed writing before the start of the experiment, this was a risk we had to accept. Although our earlier work (Wright 1990a) showed that *initially* the decomposition of letters into strokes is different for the left and right hands, we hoped that, after more practice, the g-strokes identified for the left hand would be the same as those for the right hand. This seems plausible since, as also confirmed in Wright 1990a, the overall letter shape of the letters produced is quite similar across hands.

For this analysis, we had all our subjects write ten repetitions of each letter of the alphabet in each of two contexts with their right hand. In each context, subjects wrote a connected, cursive string of letters made up of the target letter flanked by the context letter. The two contexts were *l_l*, for which the transitions to the letter are at the baseline, and *o_o*, for which the transitions to the letter are at the midline. The letters in this sample were then segmented by using velocity minima, and the strokes were described in terms of five continuous dimensions: overall curvature, angle of a line drawn between the start and endpoint, the length of the same line, and the horizontal velocity at the beginning and the end of a stroke[6]. The goal of this procedure was to characterize each of the strokes in our handwriting samples quantitatively. At this point, the data

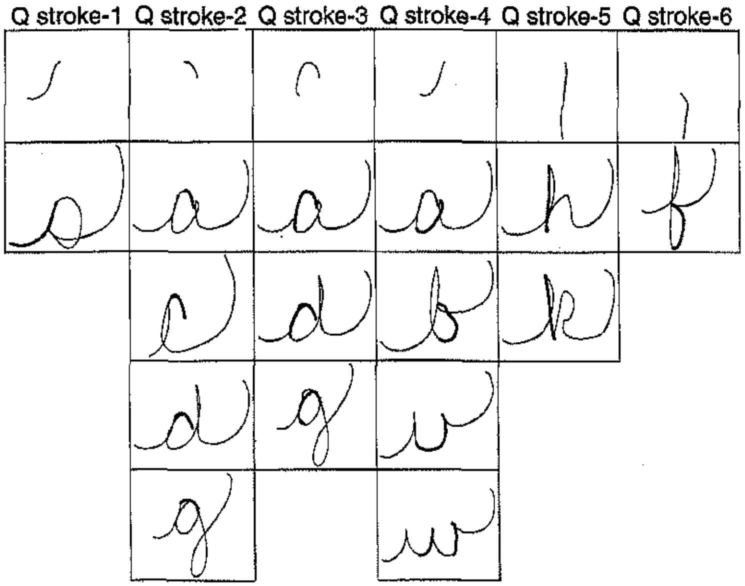

Figure 11.5
The top row shows the strokes in a letter q as identified in figure 11.2. Letters containing other examples of these generic strokes are illustrated in the panels below each stroke. The similar stroke in each of these other examples has been emboldened.

consisted of ten samples (from the repetitions) of each of slightly more than 200 stroke types (26 letters × 2 contexts × an average of about 4 strokes per letter) described on each of the five measures. These descriptions became the basis for a statistical analysis that grouped the 200 stroke types based on their similarity according to the five measures. This procedure resulted in 30–40 groups of similar strokes for each of our subjects, with each group of similar strokes corresponding to one g-stroke.

Figure 11.5 shows an example of the results of this analysis for one subject. The panels in the top row of the figure show the six g-strokes identified for the letter q (see figure 11.2) broken out into separate panels. Below each of these panels are examples of other letters containing these same g-strokes. In each of these examples, the instance of the g-stroke has been emboldened so that it stands out visually from the other strokes in the letter. Notice that the instances of the g-strokes are not always exactly the same as the examples in the panels of the second row. Although their shapes remain similar, these g-strokes are somewhat modified in the context of different letters.

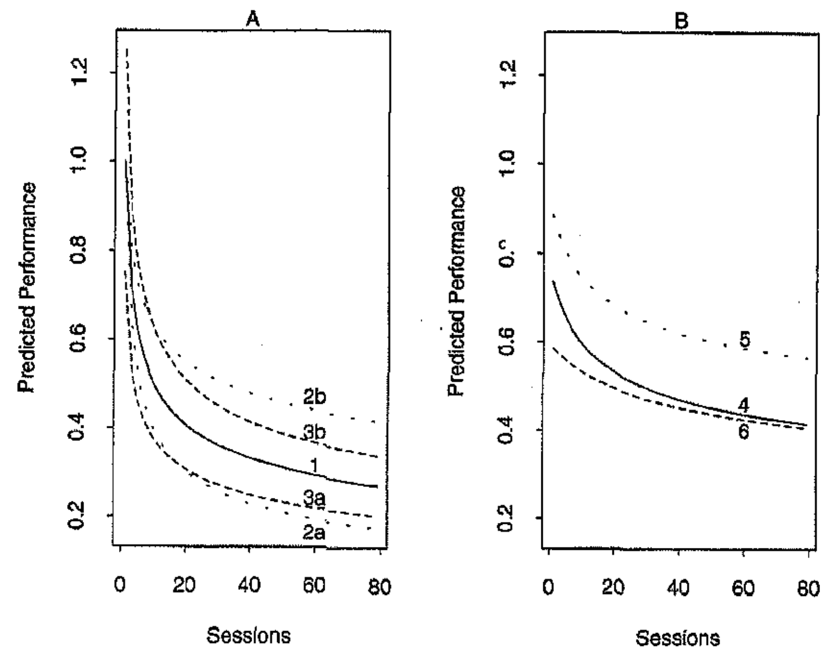

Figure 11.6
Examples of the power law learning function (equation 11.2). Panel A illustrates variations in the initial performance level p_0 and the rate of learning r. Panel B illustrates variations in the amount of prior learning n_p and the asymptotic performance level p_∞.

11.7.2 Power Law Learning Functions

Learning to write with the left hand is a complicated task that could take years to master. Curve 1 (the solid line) in panel A of figure 11.6 (which is similar to those in figure 11.4) illustrates one way that performance in this task might improve with practice. To implement the logic outlined in section 11.6, we need to be able to compare a subject's performance across changes in the words (or, as we will see below, word sets) being written. Because we cannot afford an experiment of several years' duration, we will be making these comparisons while performance is still improving from day to day. These improvements complicate the comparisons we wish to make, because they could offset small decrements in performance.

The problem of looking for discontinuities in learning that are superimposed on the overall learning trend is exacerbated by measurement error in the data: the actual data points will undoubtedly vary above and below the true learning curve. This day-to-day variation will substantially increase the size required of a change due to new materials before it is statistically discernible.

One way of reducing the severity of both of these problems is to describe the data, both before and after the change of materials, by a fitted function. We can then assess the changes in learning that occur when new words are introduced by looking for changes in the coefficients of the fitted function. One advantage of this approach is that the overall learning trend can be factored out when we look for changes in performance at the boundary between word sets. This method also makes efficient use of all the data collected across practice (as opposed to simply comparing performance immediately before and after the switch in materials, which would use only a small portion of the data). A third advantage of making the comparisons by fitting learning functions is that this method gives us a way of deciding whether the speed of learning changes after a switch in materials (recall that this was one possibility discussed in section 11.6.2).

Newell and Rosenbloom (1981) have looked extensively at how skill improves with practice for various skills. They found two characteristics of the improvements in data from different experimenters studying a variety of skill-learning domains (e.g., motor behavior, memory, problem solving). First, the improvement apparently continues indefinitely (Newell and Rosenbloom cite one data set in which there is evidence that workers continued improving at their job over more than ten years). Second, there are diminishing returns of additional practice: the rate of improvement—that is, the change in performance per unit of practice—declines as practice continues. Based on these data and their generalizations of it, Newell and Rosenbloom argue persuasively that in learning situations like ours the improvement as a function of the amount of practice can usually be characterized best by a "power law." The curves in figure 11.6, as well as those in figure 11.4, are all examples of power law learning functions.

Equation 11.1 is the simplest form of the power law learning function that applies for measures of performance, such as time, for which improvement leads to a reduction in the performance measure. In this equation, $n > 0$ represents the practice measured as sessions. The value predicted by equation 11.1, p_n, is the performance during session n. The parameter p_0 represents the level of performance on the task that a subject with no prior experience would have.[7] The parameter $r \geq 0$, in equation 11.1, is the exponent of n and quantifies the rate at which performance improves. Curve 1 in figure 11.6A was generated from equation 11.1 with $p_0 = 1$, and $r = 0.3$. Curves 2a and 2b (drawn with dots) in figure 11.6A illustrate the effect of changes in the rate-of-learning parameter, r, with p_0 held constant. A comparison of curves 1, 2a, and 2b shows that the performance measure decreases more quickly for larger values of r. Curves 3a and 3b (drawn with dashes) in figure 11.6A illustrate the effect of changes in p_0, the estimate of the initial level of performance, with $r = 0.3$. The values of p_0 correspond to the different starting points for these curves.

$$p_n = \frac{p_0}{n^r} \quad (n > 0, r > 0) \tag{11.1}$$

Because n and r in equation 11.1 are both greater than zero, with additional practice n^r grows larger and p_n becomes smaller, as expected. Notice, however, that no matter how large r becomes, p_n continues getting smaller and never reaches zero, in accord with the first of the generalizations mentioned previously. In addition, when learning follows a power law in accord with the second generalization, the improvement in performance is not constant over the sessions. Instead, the rate of improvement slows with practice. Specifically, if the first M trials are required to improve performance by the proportion D (for example, in curve 1 of figure 11.6A, the measure of performance is 1 at session 1 and 0.574 at session 4, a reduction of 42.6 percent), then to once again improve performance by the same proportion, D, will require $M(M-1)$ subsequent trials[8] (a reduction of 42.6 percent from 0.574 requires the performance measure to be 0.330, which is reached $12 = 4(4-1)$ sessions later at session 16).

Equation 11.1 can be elaborated to generalize it in two important ways. The power law functions shown in figure 11.6A all approach zero as a minimum level of the performance measure. This is called the "asymptotic level of performance." Technically, *asymptotic performance* is the hypothetical level that would be reached after an infinite amount of practice. Notice that if we substitute $n = \infty$ into equation 11.1, then the predicted level of performance is zero. Because movement time and most other measures of performance can never actually be zero, one useful elaboration of equation 11.1 is to make it possible for asymptotic performance to be something other than zero. In equation 11.2, that is the role of the new parameter p_∞. Note that as $n \to \infty$, the right side of equation 11.2 approaches p_∞.[9] A second useful elaboration to equation 11.1 helps to describe the common situation, which we noted above, in which subjects have had practice relevant to our experimental task before the start of the experiment. In such cases, the measured performance is lower than the values of p_n predicted by equation 11.1, especially for early sessions. In equation 11.2, $n_{Start} \geq 0$ estimates the effective number of sessions of practice before the start of the experiment.

$$p_n = p_\infty + \frac{p_0}{(n + n_{Start})^r} \tag{11.2}$$

Equation 11.2 is the form of the power law learning function that we will use to characterize performance as subjects practice writing with their left hand. Figure 11.6B shows sample learning curves generated by equation 11.2. As for curve 1 in figure 11.6A, $p_0 = 1$, and $r = 0.3$ for all of these curves. For curve 1 in figure 11.6A, the two new parameters in Equation 11.2 (p_∞ and n_{Start}) are implicitly equal to zero. Comparing curves 4 and 5 in figure 11.6B and curve 1 in figure 11.6A, notice how increasing p_∞ raises the height of the flat portion of a curve in the later sessions. Comparing curves 4 and 6 in figure 11.6B and curve 1 in figure 11.6A, notice how increasing n_{Start} eliminates more of the steep portion of the curve in the early sessions without appreciably changing the height of the flat portion of the curve in the later sessions.

11.7.3 A Final Design Issue

Comparing performance across word sets by fitting power law functions to the data eliminates several potential problems in evaluating the effects of a change in word set. There is, however, still one additional problem with which we must be concerned in the design and analysis of the experiment. Consider an analysis in which learning functions were fit to some chosen measure of performance. For example, we might average writing time across all of the words produced in each session. When the subject switches from one set of words to the next how much of any changes in the learning curve might be due to inherent differences in fluency between word sets?

For the analysis just described, we would not be able to make this determination unambiguously. For example, if the words in our two word sets had somewhat different numbers of strokes, then any difference we observed might simply be due to having to make more strokes in one case than in the other. We could fix this problem easily: either we could keep the average number of strokes constant across the word sets or we could first divide the writing time for each word by the number of strokes it contained before taking the average for each session (converting the observed data from the average time per word to the average time per stroke).

Eliminating differences in the stroke number across the comparisons, although a step in the right direction, is probably not sufficient to make the comparisons interpretable. An additional problem is that we do not know, a priori, whether some g-strokes are harder to produce—for example, take more time to produce—than other g-strokes, although we suspect that this is the case. Thus, if the words in the word sets to be compared are composed of different distributions of strokes, we can never be sure whether any change in performance associated with the change of word sets is due to different levels of fluency for the two sets of materials, or simply to the differences in difficulty of the strokes. (Although we found

three words—*team*, *mate*, and *dine*—to use in the examples in section 11.6 that all consist of rearrangements of exactly the same g-strokes, creating balanced sets such as these on a large scale would be difficult.)

Our solution to this problem has two parts. First, in constructing the word sets, we included all of the g-strokes in each set, although we were unable to ensure that each g-stroke occurred exactly the same number of times in each set. Thus our second step was to statistically balance the contribution of each g-stroke to the performance measure we used to summarize each session.[10]

Although these two steps eliminate a potentially serious ambiguity in interpreting our data, there is a cost to using this procedure. Because each of the word sets is constrained in this design to include all of the g-strokes, we cannot have a word set that introduces strokes the subject has never before written with the left hand. In terms of the example in section 11.6, this would be like having you learn a fourth word, which introduced new strokes, after you had finished writing *dine*—a word such as *just* or, in our earlier example, *skyscraper*. Although it clearly seems to us that it would be useful to include this manipulation, we decided that we were willing to forgo it for two reasons. First, there did not seem to be any good way to analyze the data; it would require a detailed and elaborate model to allow us to predict and thus factor out stroke differences between word sets. Second, we mistakenly expected that we would find evidence for hand specificity at the letter level and thus, given the logic we have outlined in section 11.6, a subsequent test for hand specificity at the stroke level probably would yield no new, important information.

11.7.4 Notes on Method

We have data in this experiment from three subjects for about 80 sessions of 20–30 minutes each. The subjects participated in three to five sessions per week over a span of several months; in each session, they practiced writing seven words. Each word was first written with the right hand and then repeated five times with the left hand. The right-handed writing sample remained visible throughout, and subjects were instructed to imitate their right-handed writing when writing with the left hand. This appears to be generally what they did. As each word was being written, the x- and y-positions of the tip of the pen used by the subject were digitized (recorded on a computer) with 100 pairs of x, y values recorded each second.

In this chapter, we will use the data of only one of the subjects to illustrate the common results. Table 11.3 shows the three word sets learned by this female subject. In the first twenty-seven sessions, she learned the set of words labeled "Initial Training" in table 11.3. These seven words are constructed from eleven letters, with each letter occurring at least

Table 11.3
Stimulus words used in the three phases of the experiment for the subject whose data are reported

Phase	Stimulus words						
Initial training	noon	enough	where	worth	poet	group	tower
New words, same letters	tune	report	hoop	gown	worn	huge	whether
New letters, same strokes	beam	major	quest	sobs	bombed	squad	jade

twice (the most frequent letter, *o*, occurred seven times). These eleven letters were chosen because they include instances of each of the g-strokes identified in the right-handed handwriting of this subject.

Starting at session 28, the subject stopped practicing the training set and began practicing the second set of words (labeled "New Words" in table 11.3). Practice with this set of words continued for 34 sessions. The seven words in the new word set are composed of the same eleven letters used for the training set. Thus these words differ from those in the training set only in the sequencing of the letters in the words, not in the letters or strokes used.

Finally, during the next nineteen sessions, the subject practiced using the third word set (labeled "New Letters" in table 11.3). This word set contains at least two instances of seven letters that the subject had not previously written using her left hand. (In addition, there were one or more instances of six of the eleven letters practiced previously.) Note, however, that according to the g-stroke analysis, no strokes the subject had not previously written are required to produce the seven new letters in this set.

11.7.5 Quantifying Writing Skill

We were interested in the general progress of the subjects as they practiced a particular word set and how performance changed when they switched from one word set to the next. Because improvements in writing skill are reflected in many aspects of performance, we computed measures for a number of characteristics of each stroke. Based on preliminary analyses, we chose six of these characteristics, listed in table 11.4, to describe each instance of a g-stroke. Although the angle and length measures used here are similar to those used to identify the g-strokes, the other measures are different. These differences reflect different goals. For the classification of strokes into g-strokes (see section 11.7.1), we used measures that reflect primarily stroke shape and not the manner (e.g., speed) in which a stroke was produced. Here we wish to track changes in both shape and speed measures.

Table 11.4
Stroke characteristics used to follow the improvement in writing skill over practice

Characteristic	Description
Duration	The elapsed time between the stroke endpoints.
Length	The distance between the stroke endpoints along the path of the stroke.
Peak velocity	The fastest that the pen moved during the stroke.
Angle	The angle between a line connecting the stroke endpoints and a horizontal line to the right (see figure 11.7).
Clockwise curvature	The maximum distance between the bulge of the stroke in the clockwise direction and the line joining the stroke endpoints, normalized by the length of that line (see figure 11.7).
Counterclockwise curvature	The maximum distance between the bulge of the stroke in the counterclockwise direction and the line joining the stroke endpoints, normalized by the length of that line (see figure 11.7).

11.7.5.1 Measuring Stroke Angle and Curvature

The angle and curvature measures need some clarification. Examples of the application of these measures are shown in figure 11.7 using the first four strokes in the letter *q* (see figure 11.2 for definition of the strokes). In this figure, each stroke is shown with an arrowhead marking the end of the stroke. The overall stroke direction is indicated by the angle between the solid, straight line connecting the ends of the stroke and the broken, horizontal line passing through the beginning of the stroke (with 0 degrees indicating a horizontal stroke to the right).

The amounts of clockwise and counterclockwise curvature in a stroke were measured as the maximum distances, in each direction, between the stroke and the straight line joining its ends, divided by the length of the line joining the ends of the stroke. In figure 11.7, the perpendicular segments indicating these curvature measurements are shown. The heavy, solid line marks the measurements in the counterclockwise direction, and the light, solid line marks those in the clockwise direction. (If the stroke is rotated so that the line joining the start of the stroke to its end is oriented to the right, then the perpendicular for the counterclockwise measurement goes down from this line joining the ends of the stroke, and the perpendicular for the clockwise measurement goes up.) In normal cursive handwriting, most strokes have a strong counterclockwise curvature.

11.7.5.2 Using Variability to Track Consistency

For each of the six measures listed in table 11.4, in addition to a mean value for each session, we obtained a *variability measure* as an indication of

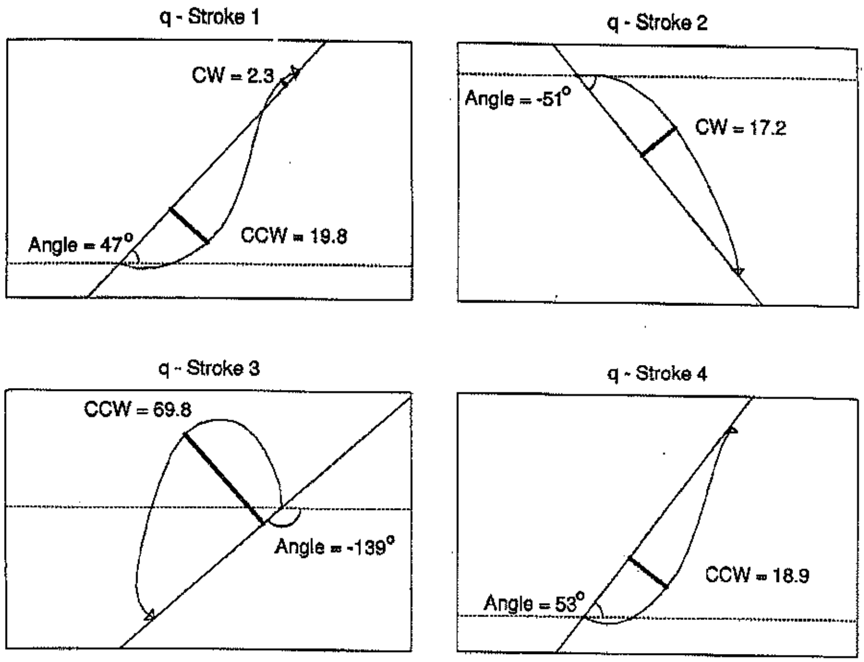

Figure 11.7
Four strokes from the letter *q* (see figure 11.2), illustrating the measurement of stroke direction and the clockwise and counterclockwise components of each stroke.

the stability of current performance. In general, a variability measure looks at the differences among a set of values, with higher values indicating a greater spread (see Wickens, chap. 12, this volume). To illustrate, in our data high variability on the size dimension would suggest inconsistency in writing size across each of the five trials per word within a session. Low variability would indicate that writing size remained fairly constant across the repetitions in a session. We obtained variability measures on each dimension for each stroke within a session. Because these variability measures reflect writing inconsistency, we expect that their values should decrease as subjects' left-handed writing becomes more fluent. There are two reasons for this expectation. First, skilled writing, being less disrupted by control problems during production, should be more consistent than unskilled writing. Second, as skill increases, the underlying motor plans should be subject to fewer reorganizations intended to improve their operation, and the shape variability associated with these changes should decrease.

Table 11.5
The five measures used to construct the composite measure of writing fluency used for the analysis in figure 11.8

Speed measures	Shape variability
Mean stroke duration	Mean variability of stroke angle
Mean variability of stroke duration	Mean variability of counterclockwise curvature
Mean stroke peak velocity	

11.7.5.3 A Composite Measure of Writing Fluency

Table 11.5 lists five of the twelve possible measures that can be created by computing the mean and variability of the six characteristics in table 11.4. We chose these measures because they provide us with representative evidence of the variety of changes that occurred with practice. Results from the other measures were either inconclusive or followed in the same pattern as those presented here. Our initial presentation of the results will be based on a composite that combines these five measures in a manner to be described and motivated below. Later, we will present the results for each of these measures individually.

11.8 Results: How One Righty Learned to Write Lefty

Given the potential complexity of the results from this experiment, the actual results are surprisingly straightforward. Figure 11.8 shows the results for the composite measure. Each point in this figure represents one session's performance averaged across the twenty-four g-strokes used to produce the seven words written with the left hand during each session. The vertical lines dividing the figure mark the points of change from one word set to the next, as labeled at the top of the figure. The units on the *y*-axis in this figure measure performance as a proportion of the initial performance, with a lower score indicating more fluent handwriting.

11.8.1 Overall Improvements with Practice

Consider first how the subject's writing improved with practice. You can get a good idea of her improvement over time by looking at the shaded curve. As you can see the most substantial improvement in performance occurred in the earlier sessions—the curve is steepest at the beginning. By the end of the study, improvement in performance leveled off considerably. This curve is a power function (see section 11.7.2), with the form shown in equation 11.2. For this fit, $p_0 = 1.8$, which is substantially larger

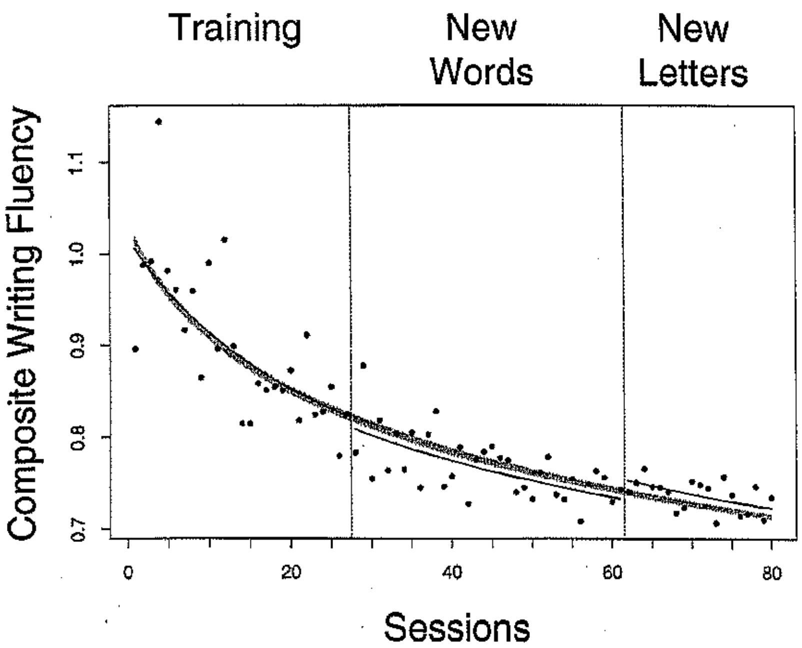

Figure 11.8
Results for one subject in the three phases of the experiment. The dependent measure is a composite measure of writing fluency. Each point represents the data from one session: the average of 5 replications for each of 24 g-strokes. The solid lines are a fitted power function with separate starting time parameters for each phase of the experiment. The shaded line is a power function fit with a single starting time parameter. The vertical lines separate the data from the three phases of the experiment.

than 1.14, the largest value observed in our data. This means that fully unpracticed performance would have been substantially worse than what we observed. Consistent with this, the fitted $n_{Start} = 16$, which means that our subject's prior experience was the equivalent of 16 sessions of practice. (This experience need not necessarily have involved handwriting per se because other fine-motor activities involving the left hand might provide preparation for the process of learning to write.) The fitted asymptote, $p_\infty = 0.4$, gives an estimate of how good performance might eventually become in this task. Finally, the rate of learning is quantified here by $r = 0.38$. Equation 11.3 shows the power law learning function with the estimated values substituted for the parameters.

$$p_n = 0.4 + \frac{1.8}{(n+16)^{0.38}} \tag{11.3}$$

At this point you may want to consider the results of your own left handed writing practice. If you have been writing *handwriting* at the end of each section of this chapter you will now have written this single word about 175 times. In figure 11.8, the results from the first 22 sessions include 110 repetitions for each of the 7 words. You can see that over the course of this practice this subject's writing improved considerably. It is likely that your own writing has also improved considerably. What kinds of improvement have you noticed? Is your writing more consistently accurate? Faster? Less jerky? As you read on about the ways we assessed improvement in our subject's handwriting, you might want to consider how your improvement might be assessed using these methods.

11.8.2 Changes in Fluency When the Words Being Written Change

Thus far we have described the overall shape of the learning function. Now we turn our attention to our central interest in these data, describing what happens when subjects switch to previously unpracticed materials. For this female subject, the first change of materials occurs after session 27. Figure 11.8 shows that there is certainly no substantial degradation of performance after this point. If there were, we would expect to see an abrupt increase in the values of the composite measure of fluency on the new materials. The data points do not seem to reflect any increase at all. Instead the data points and the shaded curve seem to continue along an uninterrupted path across this change in the materials being written; if there is any change it is downward rather than upward.

Although, looking at the data and the shaded curve in figure 11.8, this interpretation seems obvious, we would like to support it in a way that is more objective, that relies less on our individual, visual interpretation. To this end, we fitted the data in figure 11.8 with a function that is somewhat more complex than the one that we presented as equation 11.2. This function reflects changes that occurred with the introduction of new word sets. The three curve segments in figure 11.8, show the results of this fitted function. To assess whether there were changes in performance at the word set boundaries, this fitted function needs to have parameters that can reflect changes associated with word sets rather than being constant across the full range of the data. What we did was to elaborate the parameter n_{Start}, which in equation 11.2 estimates the amount of prior practice, into a trio of related parameters that we refer to as $n_{Word\,Set}$. As we have set up these parameters, their effects are cumulative across sessions. Thus the contribution of $n_{Training_1}$ starts with the first session and continues through all three word sets. It was this value that we reported as that of n_{Start} several paragraphs above. The contribution of $n_{New\,Words}$ is added beginning at the second word set and continues through the third word set. What this means is that for the first word set, the estimated number of sessions of *effective* practice, that is, the quantity raised to the power r, is $(n + n_{Training_1})$. Starting with session 28, however, the number of sessions of effective practice is estimated as $(n + n_{Training} + n_{New\ Words})$. If there is no change in fluency at the word set boundary, $n_{New\ Words}$ will be zero. If changing the words being practiced hurts performance, $n_{New\ Words}$ will be negative. And if the change for some reason improves performance, $n_{New\ Words}$ will be positive. Similarly, the parameter $n_{New\ Letters}$ estimates the *change* of prior effective practice that occurs starting with word set 3.

11.8.2.1 From Word Set 1 to Word Set 2: Same Letters, New Words

Returning to figure 11.8 and looking now at the individual curve segments corresponding to the word sets, we can see there is a small discontinuity of the fitted function at the vertical line that marks the switch from word set 1 to word set 2 after session 27. The direction and size of this discontinuity is described by $n_{New\ Words}$ for which the estimated value is 3.4 sessions. This means that changing to new words made up of the same letters may have improved performance as if the subject had done 3.4 sessions of practice outside of the experiment. Given the variability of the data (for this and our other two subjects), however, this improvement is small enough as not to be statistically distinguishable from zero. With reasonable confidence, we can only state that the true value of n_{p_2} could be within the range from 9.4 sessions of improvement in performance (i.e., additional effective practice sessions) to 2.6 sessions of degradation of performance.[11] This is not as precise an estimate of the change associated with n_{p_2} as we might wish. Notice, however, that this range excludes as unlikely the possibility that there was much degradation in performance but leaves open the possibility that there was a moderate improvement in performance.

You may find it strange that these results leave open the possibility that the subject spontaneously improved in her left-handed writing skill when she began writing word set 2. Even if the slight improvement seen here were statistically reliable, we doubt that spontaneous improvement would be the cause. It is possible, however, that the subject was more motivated to do well when we gave her new words to write, perhaps because she had become bored writing the first word set. Despite this possibility, we are comfortable drawing the conclusion that the introduction of word set 2 did not disrupt writing performance by the subject. This result is an example of outcome A in table 11.1 (see section 11.6.2). Thus we infer that the underlying representation at the word level is hand-independent and that prior learning at the letter and stroke levels is not word-context-specific.

11.8.2.2 From Word Set 2 to Word Set 3: Same Strokes, New Letters

After session 61, when there was a switch to a third word set with new letters based on the previously practiced strokes, the curve segments corresponding to word sets 2 and 3 in figure 11.8, suggest a degradation in performance. Because the power function, at this point, has flattened out, the estimate, $n_{New\ Letters}$, of what appears to be a small change in the fluency measure is equivalent to almost nine sessions of practice. Statistically, however, this difference is not distinguishable from zero. Again, we need to interpret fluency changes carefully, since sessions do not provide a precise measure of loss of fluency here. By the same standard of confidence used above, the true value of this change could range from the equivalent of two sessions improvement to twenty sessions decrement. One reason that the estimate of $n_{New\ Letters}$ is so imprecise is the flatness of the learning function in this region. Consider that, at this point in practice, even twenty sessions of practice is equivalent to a change of only 0.03 (3 percent) on the performance scale used for the ordinate in figure 11.8.

To summarize, we do not have sufficient evidence to assert that writing words containing new letters disrupts performance; neither can we assert that there is no disruption at this boundary; indeed, our best estimate of n_{p_3} indicates that a decrement in performance did occur. We can be reasonably certain, however, that there was no major disruption of the subject's ability to write with her left hand when asked to produce letters that she had never before produced. In terms of table 11.2, our results could be consistent with those of either outcome A or outcome C. We can safely exclude outcome D. Thus, based on the data from this subject alone, we cannot safely decide whether there is, in fact, hand dependence at the letter level of the underlying representations.

When the results exhibit ambiguity such as this, the most powerful tool that an experimenter has available to resolve the ambiguity is replication. In this case, we have data from two other subjects to help resolve this issue. For one of these subjects, there was no change with the introduction of words containing new letters; for the second, there was, if anything, a slight improvement in performance. Taken together, the results from these three subjects constitute strong evidence that representations at the letter level are hand-independent.

Before leaving this analysis, we should examine one of its features. Notice that we have assumed that possible changes in fluency are best characterized by changes in $n_{Word\ Set}$ rather than changes in the rate or asymptote parameters. In elaborating equation 11.2, we could have introduced triples of these other parameters. Perhaps most plausible would have been a trio of parameters $r_{Word\ Set}$, estimating different rates of learning for the different word sets. Recall that section 11.6.2 raises the possibility that there might not be direct generalization of learning but that there still could be an improvement in the learning process. This would be reflected in the learning curve by an increase in the rate of learning, represented by an increase of the parameter r. We did not include separate rate parameters for each word set in this fit because our initial analyses showed no hint of unequal learning rates. Informally, we can look at figure 11.8 and verify that there is no tendency for the data points to be above the fitted line at the start of a word set and below it at the end of the word set. We would expect to see just such a pattern, however, if the r parameter underestimated the rate of learning within one of the word sets.[12]

To summarize, although we explored adding the $n_{WordSet}$ parameters, their contribution was relatively small. The power function with a single n_{Start} parameter (the shaded curve in figure 11.8) fit the data very well. We thus conclude that the addition of additional parameters to the power function does not substantially improve our ability to describe the data.

11.8.3 Results for the Component Measures

11.8.3.1 *The Problem of the Speed-Accuracy Trade-Off*

Figure 11.8 displays data from a measure of writing fluency that is a composite of shape variability (accuracy) and speed. It is a general finding in many areas of human performance that people are adept at trading speed for accuracy (Pachella 1974). The advantage of the composite is that combining the measures for speed and accuracy roughly cancels out a strong trade-off between speed and accuracy that exists in these data. For handwriting, this means that an increase in speed tends to lead to less consistent, messier writing, while an attempt to produce high-quality, accurate writing tends to lead to slower performance. In creating the composite measure, we constructed a general speed measure and a general accuracy measure and averaged the two together.

The five measures listed in table 11.5 were combined to create this composite: three measures related to writing speed and two measures of accuracy. Before being combined, peak velocity values were inverted so that they decreased with practice, as did the other measures, and the scale of each measure was adjusted so that the speed and accuracy measures contributed equally to the overall composite.

To display their separate contributions, plots of the five measures versus sessions are shown in figure 11.9. One striking feature of these data is that there is a large discontinuity during training for each of the measures. The source of this discontinuity will be further considered below. Because of the discontinuities, the separate measures are not well fitted by power functions. Thus the fitted lines in our displays of these data are computed using local regression techniques ("loess fits"; Cleveland 1979)

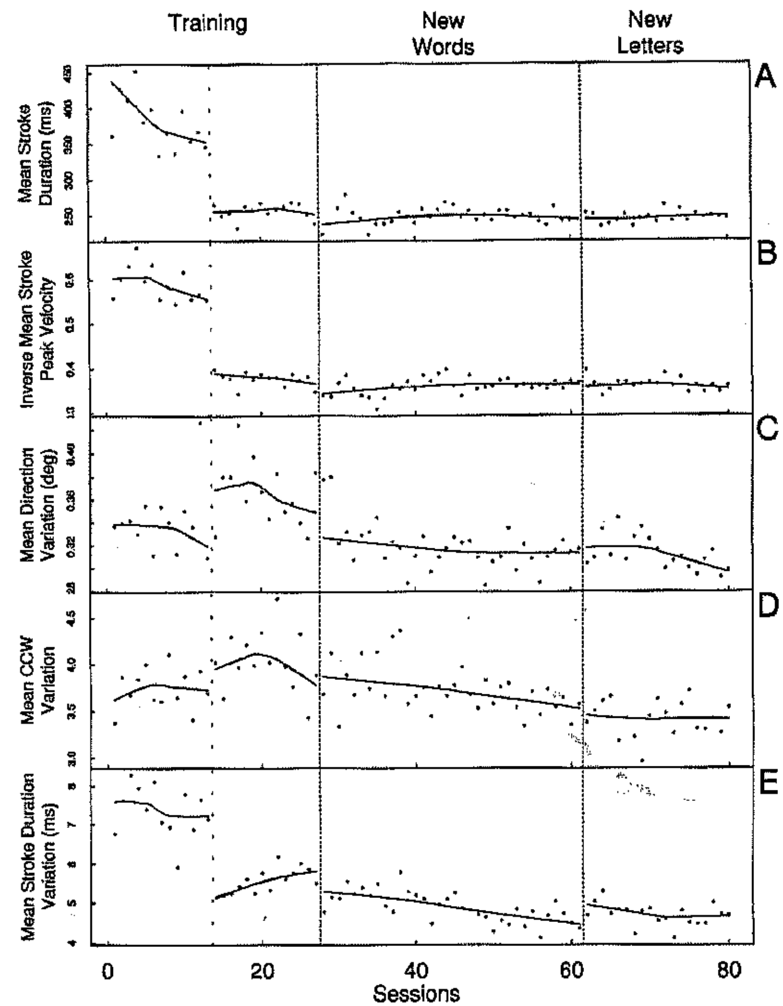

Figure 11.9
Raw data for five measures that were combined to form the composite measure of writing fluency displayed in figure 11.8. Each panel in this figure shows results for a different measure of performance. Each point represents the data from one session: the average of 5 replications for each of 24 g-strokes. The solid lines are loess functions fit separately to each section of the data. The vertical, dotted lines separate sections of the experiment. The two narrowly spaced dotted lines mark the points at which new words were introduced. The widely spaced dotted line marks the point where the subject was instructed to try to write faster.

applied separately to each word set and to each of the two parts of the first word set.[13]

11.8.3.2 Mean Stroke Duration

Figure 11.9A shows mean stroke durations as a function of practice. We expect stroke duration to decrease with practice, and the results shown here confirm this expectation. Notice, once again, that there is little if any evidence for degradation in performance at the two word set boundaries, indicated by the vertical lines.

Another important aspect of figure 11.9A is the large gap between session 13 and session 14 in the training phase marked by a dotted vertical line in the panels of figure 11.9. This gap is the direct consequence of our misguided intervention as experimenters. Because the writing speed measure we were using to monitor performance from session to session for this subject was crude, we seriously underestimated the degree to which her writing was speeding up. Thinking that her (underestimated) writing speed might simply be a failure to understand our goals for the experiment, we suggested to the subject, at the start of session 14, that it would be good if she concentrated on increasing her writing speed. Clearly, this simple instruction had a dramatic effect. As mentioned above, people can generally increase their speed on a task, but this improvement in speed tends to be at the expense of accuracy. In this case, the subject responded to our request by reducing her writing times. As we will see shortly, she achieved this dramatic improvement in speed by substantially increasing the messiness of her writing which, in our data, is reflected by increases in the shape variability measures.

Fortunately, only this subject showed such a large change in speed-accuracy trade-off. There is, however, every reason to believe that similar, if less dramatic changes, occurred in the data for all three subjects. It is because of this possibility that we prefer to use our composite measure of fluency to summarize these data as in figure 11.8. That such a simple averaging process should do so well at counteracting this trade-off is surprising. However, several attempts to come up with a better way of eliminating the speed-accuracy effects from the data have failed to locate a method that does noticeably better than this simple average.

11.8.3.3 Mean Peak Velocity

Figure 11.9B, which displays the mean stroke peak velocity across sessions, tells much the same story as the duration data in figure 11.9A. To make these data visually similar to the other measures, we have plotted the reciprocal of the peak velocities here (time per unit distance). Because peak velocity, unlike mean stroke duration, is not influenced by exactly where we place the stroke boundaries, the peak velocity data provide a useful confirmation of the pattern of results observed in the mean stroke durations. In these data, the only evidence for degradation of performance at the boundaries between word sets can be seen after the boundary at session 61, where there is a slight hint that performance may have been degraded for one session. Although the change is small and could easily be due to random variation, it must be noted that this datum represents a performance worse than that of any nearby session.

11.8.3.4 Stroke Shape Variability

We expect that the variability of stroke shape descriptors will decrease with practice. Figures 11.9C and 11.9D show how the two measures of shape variability (see table 11.5) change with practice. These data have more nonsystematic variation than the speed measures, probably because variability measures are inherently more variable. Notice also that each of these two measures shows an increase (degradation of performance) at session 14, which is when we encouraged the subject to try to write faster. Although not as clear-cut as in the speed data, this increase seems consistent with the shift in the speed-accuracy trade-off that we believe we induced by the change in instructions.[14] Most important, however, these measures do nothing to contradict the impression that there is no large degradation in performance at the boundaries between the word sets.

11.8.3.5 Stroke Duration Variability

The data for mean variability of stroke duration are shown in figure 11.9E; in some ways these data are the most complicated and interesting. The durations of movements and the variability of those durations are often observed to increase together.[15] Nonetheless, the data do not exhibit the same pattern at the word set boundaries as the mean stroke data in figure 11.9A. There are two differences worth noting. First, consistent with the hypothesis that they are related (and that duration variability is not like the shape variability measures just described), stroke duration variability decreases abruptly at session 14, just as the stroke durations do. After that, however, notice that the stroke duration variability does not remain at its new low value, but instead increases, eliminating a large part of its earlier decline. The second, and possibly more important difference, is that stroke duration variability is the one measure that shows definite evidence for degradation at the second boundary, when new letters were introduced.

11.8.4 The Development of Hierarchical Control at the Stroke Level

One question that our research design does not explicitly address is how a person learning to write might acquire the proposed hierarchical representation for handwriting control. How might units at the stroke, letter,

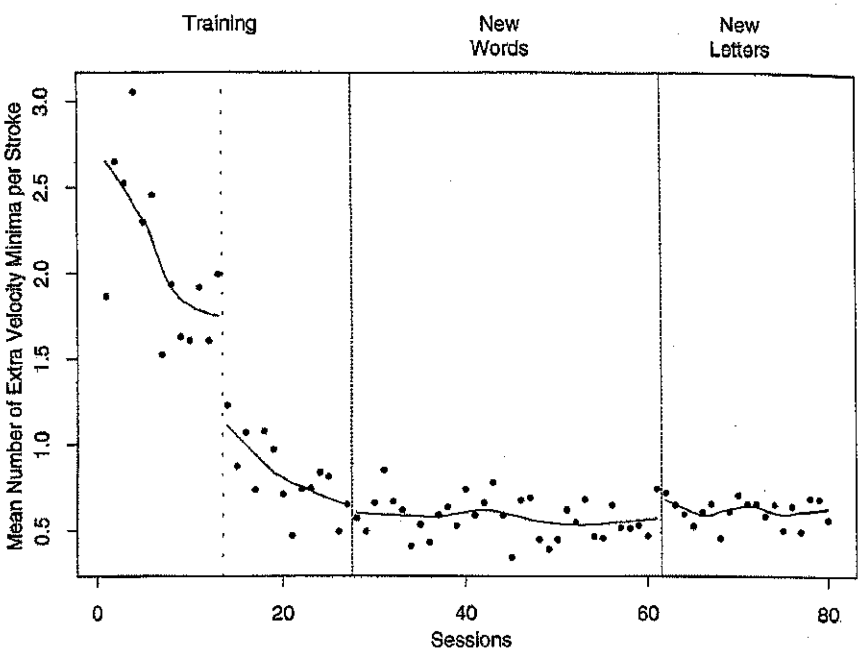

Figure 11.10
Mean number of extra tangential velocity minima per stroke across sessions. Each point represents the data from one session: the average of 5 replications for each of 24 g-strokes. The solid lines are loess functions fit separately to each section of the data. The vertical, dotted lines separate sections of the experiment. The two narrowly spaced dotted lines mark the points at which new words were introduced. The widely spaced dotted line marks the point where the subject was instructed to try to write faster.

and word level be developed? For the stroke level, at least, we have some results that partially address this issue, suggesting that there is development of higher-order units (strokes) from smaller ones (parts of strokes). The idea that a higher-level unit can be developed from lower-level units has some previous support. Pew (1966) has produced evidence supporting the idea that higher-order units are developed, in a two-key tracking task. In addition, Sternberg, Knoll, and Turock (1990) have found that skilled typists have higher-order units that encompass more than a single key press, but only in specific, limited situations. For handwriting, it is common, when a subject is first writing with the left hand, that the parts of the trajectories we have identified as strokes are not produced in a single, smooth movement (Wright 1990a). Our results show that after some practice, these same strokes are generally produced with single, smooth movements, suggesting that the strokes are being treated as individual

units. You have probably noticed this type of improvement in your own left-handed writing practice.

We can crudely quantify the change from jerky strokes made up of many segments to smooth, consolidated strokes by counting the number of extra minima in the tangential velocity that occur between the points identified as stroke boundaries. (Recall how, in section 11.4.3, we used tangential velocity minima to identify potential stroke endpoints.) The extra minima occur at places where writing slows down midstroke. The more extra minima there are, the less smooth the writing is. In ordinary fast, skillful writing the number of extra velocity minima per stroke approaches zero.

Figure 11.10 shows the average number of extra velocity minima per stroke across sessions for the subject whose data we have been reporting in this chapter. Notice how the number of additional minima decreases quickly as skill develops. This can be seen as an indication that the strokes that were identified for the right-handed writing are coming to be treated as units in the left-handed writing.

11.9 What Does It All Mean: Looking for the Writing on the Wall

Consider first the results of this experiment that are related to overall writing improvement. These results are straightforward. By several measures, there is considerable improvement of left-handed writing with extended practice: for example, mean stroke duration is reduced by 45 percent, and two measures of the variability of stroke shape are reduced by 33 percent. Thus, with practice, people can learn to produce left-handed writing that is substantially faster and more consistent than the writing initially produced by the left hand. In a sense, these results simply confirm what many people with broken or amputated limbs have already been forced to discover.

This substantial improvement in left-handed writing makes the interpretation of the rest of our results much easier. If the improvement had only been slight, there would be little reason to expect performance changes when switching to new materials. Fortunately, the large observed improvement provided ample scope to observe any deterioration in performance after a switch to new words. Thus, for example, there was plenty of room to observe intermediate and large effects such as those illustrated in figure 11.4B. Instead of these larger effects, however, what we found was little or no degradation when new materials were introduced.

Perhaps inevitably, our data are not completely clear about whether there is any deterioration in performance at the word change boundaries. Particularly at the boundary that we associated with effects at the letter level (the change to words with new letters composed of old strokes), there was the suggestion of a small, but statistically unreliable, degradation in performance. We describe this outcome as "perhaps inevitable" because it is the curse of all research that we can never confidently assert that there is absolutely no effect. Instead, the best we can ever hope for is to improve the precision with which we estimate parameters such as those for $n_{Word\,Set}$ through replication, better measurement tools, and more effective experimental design. In our data, although we found that zero degradation was consistent with our estimate of the change that resulted at the boundary between word sets 2 and 3, a small to moderate degradation of performance is also consistent with these results. In this case, we have available data from two other subjects. The data from these additional subjects substantially strengthen the argument that there is little or no degradation in performance at the boundary between words sets 2 and 3.

11.9.1 Conclusions about Hand Independence and Context Specificity of the Hierarchical Representation

We found no evidence for degradation of performance at the first boundary and evidence consistent with, at most, a small amount of degradation of performance at the second boundary. Given the interpretive framework we have outlined in tables 11.1 and 11.2 (see sections 11.6.2 and 11.6.3), these results lead to three important conclusions.

1. Substantial improvement in handwriting performance with practice reflects learning at the stroke level or lower. Practice with the left hand produced substantial improvements in performance that were not context-specific, that did not depend on the specific words or letters being written. For example, having learned to write the words in the first word set, the subject was then able to write the words in the second word set just as well as those in the first. This is important because we can assume that if the subject had not practiced word set 1, then her performance on word set 2 would have been more similar to that at the start of the experiment than what we actually observed at session 27. Because the increased fluency was maintained both when new words were introduced (with word set 2) and when new letters were introduced (with word set 3), this learning appears to be independent of both the word and the letter context. Such learning at the stroke level or lower might take the form of hand-specific information designating the muscle activation levels and timings necessary to form each stroke. Alternatively, the specification of stroke shape may also be hand-independent, in which case the hand-specific learning that is occurring must involve other, still lower level representations/processes.

2. The representation at the level of words is hand-independent, and the representation at the level of letters is largely, if not completely, so. This claim is based on the observation that words and letters never before written by our subject with her left hand were produced as fluently at the start of word sets 2 and 3 as words and letters that had been practiced with the left hand. Thus the subject's knowledge about the shapes of these previously unpracticed letters and words and how to produce them, which resulted from prior experience writing with the right hand, must have been available in a form that was useful for guiding writing with the left hand.

This second conclusion depends on our assumption that there are representations at the word and/or letter level whose learning leads to improved performance. Nothing in our results actually forces this assumption, although we find it quite plausible. One way to validate this assumption would be to show that the improvements due to practice writing a specific word with the right hand are reflected in similar changes in performance for the same word written with the left hand. This is difficult to demonstrate in college-age subjects, such as those we studied, because right-handed practice no longer results in substantial changes in performance. One way around this difficulty would be to study across-hand performance in younger children who are learning to write. Another approach would be to study college subjects learning stroke-oriented tasks similar to handwriting, such as drawing unfamiliar figures, for which the effects of learning with one hand on performance with the second can be studied directly. (Wright has begun to study acquisition of drawing expertise, but there are many technical difficulties and the results, which are still preliminary, will not be considered here.) Finally, Portier, Van Galen, and Meulenbroek (1990) provide support for the idea that there are representations at the word and/or letter level. They created strings of pseudoletters constructed by rearranging the strokes in familiar letters and found that subjects producing these strings with their right hand showed improvement with practice. The first and second conclusions lead to a third.

3. Information about how to write is organized in levels of increasing specificity. The interpretation that there was transfer of prior learning from two separate sources (learning independent of word context and letter context at the stroke level, and the hand-independent learning at the word and/or letter level) supports the hypothesis that the representations/processes that underlie motor performance in handwriting are organized in a hierarchy. The levels may be akin to a word level, a letter level, and a stroke level, although this is far less certain. The alternative is that, for example, all the information necessary to write the letter *t* or for instance, at a higher level, a word such as *team*, is stored in a single,

nondecomposable form. Instead, it appears that information about how to form complete letter shapes such as the *t* and information about how to generate the strokes that make up the *t* are stored and can be used separately.

11.9.2 Summary

At the start of the chapter, we asked what aspects of your left-handed writing would have to change in order for your writing to be considered fluent. The answer we gave then involved characteristics of the writing itself. It should be fast, smooth, and regularly formed. Now we are in a position to give a different type of answer, one in terms of the types of learning that have to occur and properties of the representations that have to develop to support fluent writing with the left hand.

Our three conclusions suggest that it is not necessary to practice all the possible combinations of letters or even all the individual letters to attain overall fluency. First of all, because this learning is not context-specific, left-handed practice on some combinations will transfer to the production of other combinations. Furthermore, it is not necessary to develop new left-hand-specific representations for words and letters written with the left hand because these representations are hand-independent. Some left-hand-specific learning must take place, however, at the stroke level. According to our results, the primary reason we are unable to produce left-handed writing fluently without practice is that we have not developed these low-level, hand-specific representations.

11.9.3 Looking beyond the Data

11.9.3.1 More Evidence about the Stroke Level

Our first conclusion, that the substantial improvement in performance with practice reflects learning at the stroke level or below, would be stronger if we could show that introducing new strokes leads to disruptions. As we discussed in section 11.7.3, we did not initially include a fourth set of words introducing new strokes for two reasons. First, we expected that this comparison would not be necessary because we expected to obtain different results for the letter level. Second, including different stroke populations in comparisons across word sets seriously complicates the interpretation of those comparisons.

After we had collected data from two of our three subjects, however, it became clear that it was important to know what happens when new strokes are introduced late in practice. Thus, for our third subject, a male, we added a fourth set of words that included this manipulation. This subject was not exposed to the full range of g-strokes during his practice of

the first three word sets. For the fourth word set, we introduced words that involved g-strokes that the subject had not yet written with his left hand. Because we have not yet fully analyzed the data from this subject, we have only a rough approximation of what the final results will be. However, our initial findings do indicate that there was a definite deterioration of performance when the subject switched to a word set containing letters that required unpracticed g-strokes. These preliminary data thus further support the conclusion that learning at the g-stroke level is hand-specific.

11.9.3.2 What Might Happen with Lots More Practice?

Although our female subject's left-handed writing is greatly improved, as shown in figure 11.8, it is still not as fluent as her writing with her right hand. Supporting this assertion are two observations: at the end of the experiment, the mean stroke duration for her left hand is slightly more than 150 percent that for her right hand. Similarly, the two stroke variation measures for her left hand are approximately 120 percent those for her right hand. Extrapolating the power function, fitted to the data in figure 11.8, suggests that the fluency of this subject's left-hand writing would be on par with that of her right-hand writing after roughly 2,000 additional sessions of practice. Clearly, we cannot give much credence to the detailed prediction from such extreme extrapolation, but it does raise a question. If this higher level of left-handed performance could be achieved, what would be the locus or loci of improvement within the hierarchical representation?

This question raises a number of complex issues that are beyond the scope of this chapter (appendix B considers one of these in more detail). We have raised this question to emphasize one danger of drawing conclusions that are too strong or general from a learning experiment such as ours that, of necessity, is only a small-scale simulation of the learning involved to master a complex task in real life. These limitations of our experimental design mean that we cannot rule out the possibility of fundamental changes of the processes and representations used to control handwriting that occurs much later in practice, for example, only after learning at the stroke level is largely complete.

11.10 Comparing Your Results to Ours

At this point you may want to look back at the changes in your writing of *handwriting* and complete your own left-handed writing experiment by trying out a few other new words. First, try writing the word *drawing* five times. It is made up of letters in the word *handwriting*—letters you have already practiced. Do you think you wrote it as well as you have been writing *handwriting*? Next try the word *mummy*. If your handwriting is similar to that of our other subjects, this word will be made up of g-strokes that you learned when you wrote *handwriting*. Of course your personal handwriting may be different and you could be using new g-strokes in this word. Did this word seem more difficult or similar? Did all of the letters seem about the same level of difficulty, or were some harder than others? Finally, try the word *question*. This word contains a number of g-strokes that you have not been practicing with your left hand.

What about your left-handed writing has improved and by what process has it improved? Our data suggest that you have learned how to produce fluently some of the g-strokes with your left hand. If you continue to practice, these motor programs will become still more fluent. If you are like our subjects, you have not had to relearn everything about skilled writing as you have been learning to write with your left hand. Some of the motor plans you already have from your skilled writing with the right hand can be transferred to your left-handed writing. Specifically, you can access information about how to sequence letters at the word level and even how to sequence g-strokes at the letter level because that information is stored in hand-independent motor programs.

Appendix A: Neurophysiological Levels of Motor Control

The neurophysiological system for motor control is composed of many separate layers. The complexity of this layering generally reflects the evolutionary position of the organism. Through evolutionary time, new layers have been added with the older ones retained largely intact. One implication of this process is that lower (more peripheral) levels of this control system can control simple operations independent of the influence of the higher levels. Generally, control information flows in a single direction, down from the cerebral cortex to the muscles at the periphery, and sensory information flows back in the other direction.

At the lowest level, it is muscles that must both generate the forces required to make planned movements and respond to changes of load. (The motor control field takes much of its terminology from engineering. *Load* refers to the collection of forces acting on the part of the body whose motion is of interest.) Muscles have properties that are similar to those of a spring. The pulling force exerted by an ideal spring increases linearly with its length according to Hooke's law: $F = kl$. Here F is the force produced, l is the length of the spring, and k is the spring constant. Muscles are like controllable springs for which the spring constant, k, is a function of the degree that the muscle is activated by the nervous system. For stiff springs or highly activated muscles, k is large, and the force produced increases quickly as the length is increased.

The springlike character of muscle is important because if we pick up an object that is unexpectedly heavy, or if we encounter an unexpected obstacle during a movement—for example, if the ink runs slowly in your ballpoint pen and the friction between it and the paper increases—one or more muscles lengthen and, because of their springlike nature, the muscles respond immediately by producing additional force. Similarly, pairs of springs connected in opposition, like opposing muscles acting on a joint—for example, the biceps and triceps are the large muscles on the top and bottom of your upper arm that act in opposition on the elbow—have an equilibrium point given a fixed level of activation. The equilibrium point is that position (in our example, the position of the elbow) at which the force exerted by the two muscles is exactly equal. Any movement of the joint away from this point stretches one muscle and allows the other to contract. This changes the force exerted by each in a way that tends to bring the joint back to the equilibrium point.

Although much of the behavior of muscles can be characterized as being like that of a spring, muscles are far from ideal springs. Ideally, for a given level of activation, the spring constant, k, of a muscle would be fixed. Instead, however, because of the way that muscles generate force and the materials of which they are made, their spring constant also changes with their length and with the speed with which the muscle is contracting or being stretched. For example, a muscle being stretched exerts more force at a given muscle length than a muscle whose length is not changing, and a contracting muscle exerts somewhat less force at a given length. Similarly, if a muscle is contracting quickly, it will exert less force than if it is contracting more slowly. Finally, for differing levels of steady state activation, equal changes in activation do not produce equal changes in the muscle's spring constant.

These complex properties of muscle are important to motor control researchers in that they make it much more difficult to make inferences about how a movement is being controlled. Because of these muscle properties, changes in force exerted at the hand cannot simply be ascribed to changes in the underlying control signals from a motor program. In particular, if the same movement were made with a different load (e.g., a pencil rather than a pen), the forces produced would need to be quite different. Some of these differences are due to changes in muscle activation produced by the underlying motor program, and some are due to force changes produced by the muscles themselves. It is quite difficult, however, to determine which is which.

The muscles are only one of several levels in the motor system that complicate our picture of how movements are controlled. Consider how your leg jerks upward when the doctor taps lightly on a tendon in your knee. The tap of the tendon stretches it and the attached muscle. This stretch is detected by receptors, which send nerve signals up your spinal cord and beyond. The nerve signals trigger a mechanism in the spinal cord called the "stretch reflex" that reacts to the signals and briefly increases the activation of a muscle in your leg whose contraction opposes the stretch. All of this happens within 25–40 milliseconds and without direct intervention from in the brain.

The stretch reflex is one of many, in some ways "primitive," mechanisms controlled peripherally in the motor system; that is, in the spinal cord or the brain stem. These mechanisms can produce behavior with complex features that extend in time. For example a frog whose nervous system is severed between the spinal cord and the brain will still swipe with the appropriate foot at a location on its body that is irritated by a pin prick. Even more impressive is the observation that a cat, after a similar operation, will locomote with different gaits at different speeds if its weight is supported with its feet on a moving treadmill.

Because of these and other mechanisms, there are myriad potential sources of movement modification and control interposed between a *plan* and the resulting *movement*. Although peripheral layers of the motor system may make high-level control feasible, from the perspective of studying the higher-level processes, these peripheral processes are a great source of complexity and confusion. In effect, they interpose poorly understood layers of buffering and filtering between us, as observers, and the central processes that plan or control movement sequences. Of course, this situation is hardly unique to the study of motor control or cognitive psychology in general.

Appendix B: Could Extended Training Produce Hand-Dependent Representations at the Letter and/or Word Levels?

Section 11.9.3.2 pointed out that, at the end of this experiment, performance with the left hand, although much improved, was still substantially worse than performance with the right hand. Although we can only speculate about whether right-handers can ever write as well with their left hand as with their right, we think it is likely that our subjects, left-handed writing would continue to improve, although much more slowly, with additional practice. Section 11.9.3.2 raised the question of what locus or loci additional learning would have. One possibility, which was unlikely in the context of our experiment but which we must consider again if subjects receive substantially more practice, is that some portion of the learning might take place through improvements at the letter or, possibly, the word level and not simply at the stroke level.

Taking the letter level as an example, one way that this learning could take place involves the writer collapsing all of the generic, stroke-level information needed to produce a specific letter into a new, letter-level unit. Subsequent production of this new letter unit would bypass the generic representations at the stroke level. Because the stroke movements in this new letter unit are specific to a particular letter, the procedures used to produce the strokes can be altered for this context—so that, for example, they blend more efficiently with the strokes immediately before and after—improving the overall performance for that letter. It seems unlikely, however, that these efficiency advantages come without some cost. For example, because our results suggest that stroke-level information is hand-specific, we would expect these new units also to be hand-specific. In principle, the same process of collapsing lower-level units could work at the word level, creating representations that bypass the generic ones at the letter and/or the stroke levels.

To explain why the experiment reported in this chapter provided little if any evidence for learning at the letter or word levels in the hierarchy, it seems possible that a certain degree of proficiency at the stroke level would have to be achieved before improvements at the letter or word level could take place. However, perhaps once skill is developed at the stroke level, g-strokes could then be incorporated into letter level units.

There are other places we might look for confirmation that such mechanisms are possible. Wright (1990a, 1993) provides support for the idea that there are hand-specific units at the letter/word levels. These studies showed that many subjects used special forms of certain letters when writing their first name with their right hand. Other subjects, however, did not use these special "signature" forms when they wrote their name with their right arm or left hand; instead, they used the normal letter forms.

As we suggested at the start of this appendix, it seems possible that a mechanism (perhaps similar to the one involved in the consolidation of strokes discussed in section 11.8.4) might come into play later in learning at the letter level or even, as the examples of subjects writing their names suggest at the word level. Given practical limitations on experimentation, evidence for such hand-specific, higher-level units will probably have to come from experiments that do not rely on manipulating practice.

If hand-specific, higher-level units are found, this result would not come as a complete surprise to us. It was with the expectation of finding evidence for hand-specific units at the letter level that we began the experiment reported in this chapter; we found that if there are such hand-specific, high-level units, they develop only after substantial practice and the achievement of considerable fluency.

Suggestions for Further Reading

Two excellent introductory books are Rosenbaum 1991, on human motor control in general, and Singley and Anderson 1989, on modeling transfer of cognitive skills.

Viviani and Cenzato (1985) illustrate the fascinating regularities in the way strokes are produced in handwriting and drawing, and compare these to regularities observed in typewriting, while Hayes and Flower (1980) and Flower and Hayes (1980) explore planning in a nonmotoric task; the composition of writing.

The issue of how to demonstrate that there is planning of movement sequences is an important one. Along with Wright 1990a, cited in the text, the reader would be well advised to study the original papers of Sternberg and colleagues (Sternberg, Monsell, et al. 1978; Sternberg, Wright, et al. 1982; and Sternberg, Knoll, et al. 1989).

Highly skilled individuals are often thought to possess an innate talent. Ericsson and Charness (1994) argue strongly, however, that the only difference between highly skilled individuals and others is cumulative, directed practice. Their thesis is well presented and intriguing, if somewhat controversial.

Questions for Further Thought

11.1 While reading this chapter, you practiced writing *handwriting* many times. Determine some aspect of the strings you wrote that you can measure easily using, for example, a ruler or a protractor. Make this measurement for each of your productions of *handwriting*. Divide these measured values into groups of about ten and for each group compute the mean,

$$\text{mean} = \frac{1}{10}\sum x_i,$$

and the standard deviation,

$$\text{sd} = \sqrt{\frac{1}{10}(x_i - mean)}.$$

Graph these values versus practice and characterize your results.

11.2 Think of another activity for which performance can be easily measured and for which you expect skill to increase with practice. This need not involve movement; many of the general concepts about skill learning presented in this chapter apply equally well in other areas of human activity. Design an experiment to study either the learning or the generalization of this skill with practice.

11.3 A hypothesized hierarchical representation played a key role in this chapter. What are the advantages of such a representation? What might be its disadvantages? Try to devise an experiment that demonstrates that there are hierarchical mental representations.

11.4 Watch the pen tip carefully as someone writes. You should be able to observe systematic changes in velocity. These velocity changes have been interpreted as markers of stroke endpoints. What else can you observe that fits with this conception of partitioning letters into strokes? Why is this a useful way to think about the production of handwriting? What other activities can you think of that have a similarly segmented structure?

11.5 Identify an activity that requires planning (it need not involve movement). How could you study the planning process for this activity? Do you imagine that you will see systematic differences in the way that novices and experts at this task go about the planning process? Design an experiment to study these differences.

Notes

This work was supported by National Science Foundation grant BNS-91-11085. The authors would like to acknowledge their many helpful discussions with Rebecca States during the course of their research. Nina Macdonald, Don Scarborough, and Saul Sternberg provided extensive, useful comments on earlier drafts of this chapter.

Charles E. Wright is currently at the University of California at Irvine. The authors can be reached by E-mail at pgl@psych.columbia.edu and at cewright@uci.edu, respectively.

1. An acquired agraphia, such as this one, is a decline in writing performance that cannot be accounted for by sensory or motor problems (e.g., blindness or paralysis). In apraxic

agraphia spelling is intact, but letters are not formed properly. Often, as in this case, the patient is aware of the errors in letter formation as they occur.

2. In Hollerbach's model (1981), the representation of each stroke consists of frequency and phase information describing the action of abstract oscillators acting in horizontal and vertical orientations. Using this stroke-level mechanism, there would, presumably, need to be a still lower, fourth, level of the hierarchical representation in which the controllable values for the system of oscillators are mapped onto *synergies*—temporary, functional organizations of the muscles and joints within a particular effector system—that produce, in this case, stroke movements (cf. Saltzman and Kelso 1987).
3. In outcome B, learning at the letter and/or stroke level cannot be used in a new context. This means, for example, that practice writing the letter *e* in *team* with your left hand would not enable you to write the letter *e* in *mate*. But because we have no reason to believe that the basic organization of writing processes differs for the left and right hand, this supposition should apply equally well to right-handed writing. Thus having learned to write the letter *e*, or any letter of the alphabet, in one word should not enable you to produce that letter in any other word. Without generalization of lower-level learning across contexts, it is difficult to see what learning at these levels could consist of. The only alternative we can imagine is that the hierarchy effectively consists of the word level alone. If so, what does it mean to answer yes to the question of hand independence at the word level, as in outcome B? If the word level is the basic unit, and it is effector-independent, then we would expect the word level to generalize from the right hand to the left hand. Any word that you have written with your right hand you should be able to produce with the left hand *in the same manner*. The results of Wright (1990a; as discussed in section 11.3) demonstrate that this is not the case.
4. It is also possible that the learning in either of these cases might be faster than the original learning at the equivalent level of performance. For simplicity, we have not tried to represent this alternative in figure 11.4B.
5. You might think that context-independence at the stroke level should be an issue because we assumed that outcome A holds for the transition from *team* to *mate*. However, that result does not necessarily require that there even be a stroke level. Everything that was learned below the word level during practice of *team* could have been localized at the letter level.
6. Horizontal velocity at either end of a stroke tells a lot about the shape of the stroke. Consider the horizontal velocity at the end of the first upward stroke in each of the letters *n*, *i*, and *e*. All three strokes end at the midline: for an *n*, the horizontal velocity is positive (indicating rightward movement); the cusp of the *i* is formed by keeping the horizontal velocity near zero; in an *e*, the horizontal velocity is negative (indicating leftward movement).
7. Note that, although possible, it is generally not a good idea to estimate p_0 as the measured level of performance for the first session. There are two reasons for this. First, we can never be certain that the subject has not done a task sufficiently similar to the experimental task prior to the start of the experiment. A second reason is that performance measurements for any single session in this task may not be very precise because of measurement error (see Wickens, chap. 12, this volume).
8. This assertion is not hard to prove. According to equation 11.1 performance at the first trial is given as $p_1 = p_0/1^r = p_0$. A premise of the assertion is that performance at trial M is $p_M = Dp_1 = Dp_0$. Another way of expressing performance at trial M comes from equation 11.1: $p_M = p_0/M^r$. Equating these two expressions and solving for D gives $D = 1/M^r$. To verify the assertion, we must determine the trial, which we will refer to as x, at which performance is $p_x = Dp_M = D^2p_0$. Another way of expressing performance at trial x comes from equation 11.1: $p_x = p_0/x^r$. Equating these last two expressions, substituting the value of D determined above, and solving for x gives $x = M^2$. Thus, as asserted, the number of additional trials required is $x - M = M(M - 1)$.

9. If we were to use equation 11.1 to describe data for which the true asymptote is greater than zero, the fitted values would tend to systematically underestimate the observed data toward the end of practice.
10. To balance the contributions of the g-strokes in the data analysis, we first computed the mean of the performance measure for each of the g-strokes across all the occurrences of that g-stroke within each session. We then averaged these means for the g-strokes together to obtain an overall measure of performance for the session. Because each g-stroke type contributed equally to these averages, we could compare these averages across sessions in different word sets or fit the learning curves to them.
11. This range of possible values is called a "confidence interval" in statistics. Its calculation is based on the variation of the data points in figure 11.7 about the fitted function. In this case, "reasonable confidence" means that we can expect 95 percent of statements, such as the range above, to include the true value of the mean. Conversely, however, we must realize that as often as one time in twenty, the true value will not be in the stated range (see Wickens, chap. 12, this volume).
12. It is important to realize that increasing the number of rate parameters does not eliminate the need for three $n_{Word\,Set}$ parameters because changes to the rate parameter change the shape of the fitted curve (see figure 11.6A) but cannot explain an abrupt change in fluency at a word set boundary. Including additional rate parameters adds to the total number of parameters being fitted and thus has the undesirable effect of making the estimates of other parameters less precise. (See Massaro, chap. 8, and Dosher, chap. 10, this volume, for more on the problems of parameter estimation.)
13. The local regression technique is a way to obtain a fitted function by smoothing the data. Essentially each fitted value is obtained by fitting a line to a range of data values on either side of the fitted point. In using a technique such as this we are attempting to isolate the changes across the data points that are systematic while reducing the influence of the sampling variability affecting each observation. This fitting technique is particularly useful in a situation such as this because it works without our having to specify a function, such as a power law.
14. In the direction variability data, session 11 has the highest value observed, and session 14 has a value more like those immediately preceding than those following the change of instructions. For the counterclockwise variation, although there is an apparent increase in session 14, the value for session 15 is back down at the level of those prior to session 14. Despite these lacunae, which may reflect changing strategies by the subject or simply the high variability in these data, we feel that there is good evidence for a shift in the overall level of variability measures at this boundary.
15. For our data, if we correlate stroke durations and stroke duration variability across sessions, the correlation is $r = 0.91$. (See Wickens, chap. 12, this volume, for a discussion of this measure of correlation.)

References

Castiello, U., and Stelmach, G. E. (1993). Generalized representation of handwriting: Evidence for hand independence. *Acta Psychologica* 82, 53–68.

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association* 74, 829–836.

Edelman, S., and Flash, T. (1987). A model of handwriting. *Biological Cybernetics* 57, 25–36.

Ericsson, K. A., and Charness, N. (1994). Expert performance: Its structure and acquisition. *American Psychologist* 49, 725–747.

Flower, L., and Hayes, J. R. (1980). The dynamics of composing: Making plans and juggling constraints. In L. W. Gregg and E. R. Steinberg (Eds.), *Cognitive processes in writing*, 31–50. Hillsdale, NJ: Erlbaum.

Greene, P. H. (1972). Problems of organization of motor systems. In R. Rosen and F. M. Snell, (Eds.), *Progress in theoretical biology* 2, 303–338. New York: Academic Press.

Hayes, J. R., and Flower, L. (1980). Identifying the organization of writing processes. In L. W. Gregg and E. R. Steinberg (Eds.), *Cognitive processes in writing*, 3–30. Hillsdale, NJ: Erlbaum.

Hollerbach, J. M. (1981). An oscillation theory of handwriting. *Biological Cybernetics* 39, 139–156.

Keele, S. W. (1981). Behavioral analysis of movement. In V. B. Brooks (Ed.), *Handbook of physiology: The nervous system*. Vol. 2, *Motor control*, 1391–1414. Baltimore: American Physiological Society.

Lashley, K. S. (1942). The problem of cerebral organization in vision. *Biological Symposia* 7, 301–322.

Margolin, D. I., and Wing, A. M. (1983). Agraphia and micrographia: Clinical manifestations of motor programming and performance disorders. *Acta Psychologica* 54, 263–283.

Merton, P. A. (1972). How we control the contraction of our muscles. *Scientific American* 226, 30–37.

Newell, A., and Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition*, 1–55. Hillsdale, NJ: Erlbaum.

Pachella, R. G. (1974). The interpretation of reaction time in information-processing research. In B. H. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition*, 41–82. Hillsdale, NJ: Erlbaum.

Pew, R. W. (1966). Acquisition of hierarchical control over the temporal organization of a skill. *Journal of Experimental Psychology* 71, 764–771.

Portier, S. J., Van Galen, G. P., and Meulenbroek, R. G. (1990). Practice and the dynamics of handwriting performance: Evidence for a shift of motor programming load. *Journal of Motor Behavior* 22, 474–492.

Raibert, M. H. (1977). *Motor control and learning by the state space model*. Technical report AI-M-351. NTIS no. AD-AO26-960. Cambridge MA: Massachusetts Institute of Technology.

Rosenbaum, D. A. (1985). Motor programming: A review and scheduling theory. In H. Heuer, U. Kleinbeck, and K. H. Schmidt, (Eds.), *Motor behavior: Programming, control, and acquisition*, 1–33. Berlin: Springer.

Rosenbaum, D. A. (1991). *Human motor behavior*. New York: Academic Press.

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

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

Singley, M. K., and Anderson, J. R. (1989). *The transfer of cognitive skill*. Cambridge, MA: Harvard University Press.

Sternberg, S., Knoll, R. L., Monsell, S., and Wright, C. E. (1989). Motor programs and hierarchical organization in the control of rapid speech. *Phonetica* 45, 175–197.

Sternberg, S., Knoll, R. L., and Turock, D. L. (1990). Hierarchical control in the execution of action sequences: Tests of two invariance properties. In Jeannerod, M. (Ed.), *Attention and performance XIII: Motor representation and control*, 3–55. Hillsdale, NJ: Erlbaum.

Sternberg, S., Monsell, S., Knoll, R. L., and Wright, C. E. (1978). The latency and duration of rapid movement sequences: Comparisons of speech and typing. In G. E. Stelmach (Ed.), *Information processing in motor control and learning*, 117–152. New York: Academic Press. Reprinted in R. A. Cole (Ed.), *Perception and production of fluent speech*, 469–505. Hillsdale, NJ: Erlbaum, 1980.

Sternberg, S., Wright, C. E., Knoll, R. L., and Monsell, S. (1982). Motor programs in rapid speech: Additional evidence. In R. A. Cole (Ed.), *Perception and production of fluent speech*, 507–534. Hillsdale, NJ: Erlbaum.

Viviani, P., and Cenzato, M. (1985). Segmentation and coupling in complex movements. *Journal of Experimental Psychology: Human Perception and Performance* 11, 828–845.

Viviani, P., and Terzuolo, C. (1980). Space-time invariance in learned motor skills. In G. E. Stelmach and J. Requin (Ed.), *Tutorials in motor behavior*, 525–533. Amsterdam: North-Holland.

Wright, C. E. (1990a). Generalized motor programs: Reevaluating claims of hand independence. In M. Jeannerod (Ed.), *Attention and performance XIII: Motor representation and control*, 294–320. Hillsdale, NJ: Erlbaum.

Wright, C. E. (1990b). Controlling sequential motor activity. In Osherson, D., Kosslyn, S. M., and Hollerbach, J. M. (Eds.), *Invitation to cognitive science*. Vol. 2, *Visual cognition and action*, 285–316. Cambridge, MA: MIT Press.

Wright, C. E. (1993). Evaluating the special role of time in the control of writing. *Acta Psychologica* 82, 5–52.

About the Authors

Patricia G. Lindemann is a doctoral student in the Department of Psychology at Columbia University. She writes:

> At Wesleyan University I majored in cognitive science with a focus on psychology. Through my course work in psychology, philosophy, and computer science, I learned the value of interdisciplinary insights into the workings of the mind. I also discovered that psychological research held a strong fascination. I planned to go to graduate school to study cognitive psychology after graduation, but instead I married, and my husband and I went to Morocco, teaching English as Peace Corps volunteers. It was during this time that I discovered my love for teaching.
>
> I finally began graduate work at Columbia University in 1991, with Charles E. Wright as my advisor. As a way of combining his interest in complex human motor control and my interest in human learning, we began a project that involved teaching people to write with their nondominant hands and observing how learning progressed. This project, which began as my master's research, is the basis for the results presented in this chapter.
>
> My doctoral research is more closely related to work I did as an undergraduate, focusing on decision making processes in the context of students deciding which college to attend. But questions in motor skill learning continue to capture my imagination, particularly from a developmental perspective. This interest is encouraged by my eighteen-month-old daughter, who constantly reminds me just how complicated it is to master each and every motor sequence—be it reaching, rolling over, running, or (one day) learning to write her name.

Charles E. Wright is a member of the Department of Cognitive Sciences at the University of California at Irvine. He writes:

> In high school I expected to be a biologist or physicist, but my studies at Wesleyan University were in the College of Letters—essentially a Great Books program in which my interests were philosophy of mind, history, and systems of classification. I also worked in the computer center throughout college.

I went to graduate school in psychology at the University of Michigan but left after two years to work as an assistant to Dr. Saul Sternberg in what was for me the mecca for experimental psychology, the Human Information-Processing Research Department at Bell Laboratories. Here Dr. Sternberg encouraged me to study human motor control.

Ten years later, having completed my Ph.D. long-distance, I joined the faculty at Columbia University. I had studied speech, typing, and simple movement at Bell Laboratories, but I wanted to study more complex movements. I began research on handwriting, and some of that research is reported here.

I am now at the University of California at Irvine, where I plan to extend my analysis of skilled movements to questions of how children learn motor skills. I am also developing clinical applications of my work, in particular, analyzing focal dystonia of the hand ("writer's cramp") and assessing treatments for this disorder.

I enjoy research in which new insights and new technology expand the limits of the questions we can ask, such as the research reported here. We had to develop many new techniques, such as the characterization of pen strokes from their trajectories, the analysis of "g-strokes," and techniques to manage over 100 megabytes of data. In the end, the results surprised us. I prefer this outcome because it is the unexpected that teaches us new things and forces us to reformulate our understanding of the world.