PERCEPTUAL LEARNING

Felice Bedford

I. Introduction

Most psychologists' familiarity with perceptual learning comes from the classic example of plasticity known as *prism adaptation*. That's the one instructors of all levels demonstrate to their students on a slow day. A volunteer straps on a pair of goggles, which has been fitted with a wedge prism, and tries to walk around the room. The task is not easy because the prism distorts vision so that objects are not where they appear to be. The prism displaces the visual image of the world to the side of the true location. There is something of a slapstick appeal, if nothing else, in watching the student bump into walls and knock things off desks or, when I was a student, miss the chalkboard erasers thrown for him to catch. Yet the comedy has its reward when the volunteer not only improves, but, if testing is done right, continues to have altered visual-motor behavior even when the goggles are removed. This usually comes as a surprise to the class, and to the perplexed volunteer as well.

Despite its intuitive appeal, and a long history of experimentation, prism adaptation has remained something of an anomaly both from the standpoint of perception and from the standpoint of learning. It has earned the status of *plasticity*, the label reserved for things unclassified that seem to be both learning and not learning at the same time. Research interest waned by the 1970s, stayed low in the 1980s and, so far, in the 1990s. I would like to convince readers of this article that perceptual learning, of

THE PSYCHOLOGY OF LEARNING
AND MOTIVATION, VOL. 30

which prism adaptation is a subset, should be a thriving enterprise. It is a perfect candidate for a specialized acquisition device, with its own rules, that could even serve as a model of study. While others have suggested the potential uniqueness of perceptual learning (e.g., Epstein, 1967; Gibson, 1969), they did so without the benefit of theoretical interest in modules and domain-specific acquisition systems developed largely in the last decade (e.g., Fodor, 1983; Keil, 1992; Rozin & Schull, 1988).

This article will illustrate the fruitfulness of the approach, largely with experiments from three paradigms used in my laboratory. The first consists of variations on the classic prism adaptation paradigm, designed to ask questions different from those typically addressed. The second looks at mappings with more modern advances of a computer monitor and a "mouse-like" pen and tablet, a paradigm that also adds a cognitive flavor. The third involves some new ideas about the phenomenon known as the McCollough effect (McCollough, 1965), which involves forming connections between properties of orientation and color. I begin first with some general considerations on learning.

II. Classification of Learning Processes

Perceptual learning phenomena would seem less anomalous if they were located within a broad classification of learning in general. Such categorization schemes, with or without perceptual learning, are too infrequent. One such categorization is briefly suggested here. At the top level, we begin most generally with *effects of experience*. The first division takes its inspiration from Paul Rozin and Jonathon Schull (1988), who in turn quote from ethologist K. Z. Lorenz (1981). To paraphrase, he asks us to contemplate why it is that learning always makes us better. If this philosophical comment causes hesitation, I suggest that is because it is so fundamental and integral to what is meant by learning as to go unnoticed. Like all things that are so obviously true, they get left out of definitions, theories, and classifications. An analogy may be the way three-dimensional space is taken for granted (Shepard, 1991). Shepard notes that people say things like "If only I had a larger office, I would have more room for my books"; but not "If only I had a four-dimensional office, I would have so many more degrees of freedom for arranging them!" (p. 3).

Figure 1 illustrates the division of *Experience* into "Experience that makes us better" and "Experience that makes us worse." The former category is what should be meant by learning. The latter category consists of experiences that leave us worse for wear, including the prototypic nonlearning effects such as fatigue and injury.

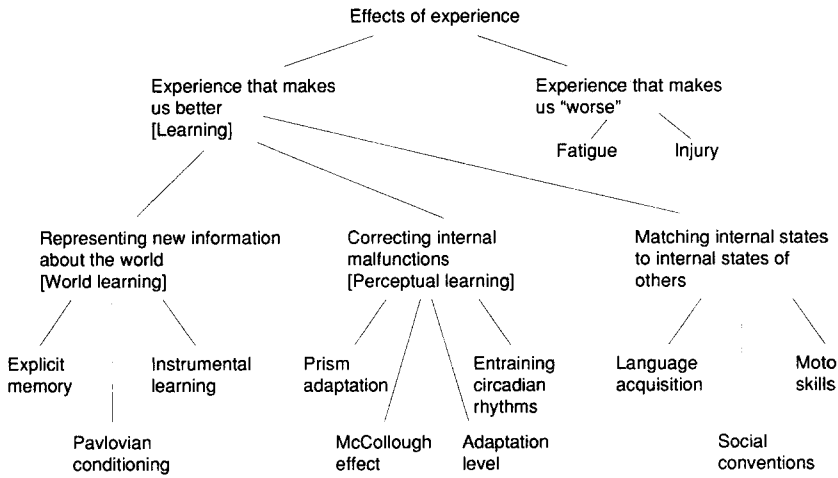

Fig. 1. Classification of different types of learning.

An illustration of how the division may capture intuition concerns the phenomenon known as *habituation*. A characterized scenario is about the family that lives by the train station, where noisy trains pass through once an hour. In reply to guests who ask how they can tolerate the trains, they remark, "What trains?" A more scientific example comes from animal behavior. An unexpected loud noise causes rats to startle; the response decreases with repetition of the noise, until the rat's behavior is uninterrupted by the noise. Not all scholars view habituation as a legitimate example of learning (see Rescorla & Holland, 1976, for discussion). The basic source of disagreement may be a matter of individual intuition about whether habituation makes us better equipped, or worse off, at dealing with the world than we were before the experience. Those who do not regard habituation as learning also tend to view the decreased responding as an unavoidable consequence of any physical system that is subjected to repeated use. On the other hand, those who think that habituation serves an adaptive function do regard the phenomenon as learning. The gist of many of these views is that it is sensible to ignore a repeating stimulus that is of little consequence, in order to free resources for other potentially more important events. The contrast suggests that the distinction in the learning classification does capture intuition about what should be regarded as learning, and as such may be useful.

Experiences that make us better can next be broken down into two major categories: processes that apprehend new information about the world and processes that correct internal malfunctions, or otherwise im-

prove sensory systems. It is the former category that is most often regarded as *learning*. Well-known examples are the associative processes. For instance, Pavlovian conditioning may be "a primary means by which the organism represents the structure of its world" through learning relations among events (Rescorla, 1988, p. 152). A distinct process also in this category is *explicit memory,* where there is storage and conscious recollection of specific learning episodes or recently presented information.[1]

Yet the other category, while less familiar, should properly be regarded as learning in the broad sense defined above. Processes within both categories make us better equipped to deal with the world. The former, which we can call *world learning* or "external learning" does so directly, through the representation of environmental properties (e.g., Gallistel, 1990, Chap. 2). The latter, *perceptual learning* or "internal learning" does so indirectly, through improvements to the organism, which in turn make us better able subsequently to apprehend the world. The perceptual learning processes keep the sensory systems in good working order to allow subsequent world learning to occur.

Before turning to perceptual learning, I will mention that there is likely also a third category. The first category, *world learning,* involves matching one's own internal states to the world. The second category, *perceptual learning,* usually involves matching one's own internal states to one another. The third category involves matching one's own internal states to the internal states of others. One example in this category is language acquisition. Languages are not out in the world (e.g., Bloom, 1990; Chomsky, 1986), nor is their acquisition a matter of correcting malfunctions. Instead, what matters is that your language match some one else's language. Another example may be the acquisition of social conventions. The function of learning that Americans drive on the right side of the road is to avoid accidents with other Americans, who also drive on the right side of the road. More tentatively, but arguably, motor skills belong primarily in this category. Skiing is not something out in the world to be acquired, but is arbitrary; its acquisition is an attempt to acquire what someone else can do. Many of the motor abilities we need in order to interact with the world mature early (e.g., walking and running); motor skills may be a

[1] *Implicit memory* (Schacter, 1987; 1992) is intentionally omitted as a type of learning. The problem with the category is not that it is heterogeneous, but that it is held together by a concept that does not make it a meaningful category. The category is defined by processes in which conscious or intentional storage or recollection of information is *not* present. In a manner analogous to the way that "not the liver" or "not tables" would not be meaningful categories of body organs or furniture respectively, "not explicit memory" isn't a meaningful category of learning. Of course, the different phenomena from which the data are collected, most frequently certain "priming" effects, still need to placed within this categorization.

matter of putting together components in the correct order. (See Adams, 1987, for a review.) Similar arguments about skills not being in the world may extend to cognitive skills, like chess playing, and other skills such as music composition.[2]

III. Perceptual Learning: Beginning Anew

There are three primary questions to be asked of any potentially distinct acquisition system (Rescorla, 1988). "What are the circumstances which produce learning? What is the nature of that learning? How is the learning manifested in behavior?" (p. 151). To this can be added a fourth: What purpose does it serve? (e.g., Rozin & Schull, 1988). Deciding when we have a "natural-kind" category is a formidable task. A separate system should distinguish itself from other systems on at least some answers to the above questions. The first three are questions about process, dealing with input, internal states, and output, respectively. The last is a question of evolutionary function. I suggest that perceptual learning distinguishes itself from world learning on all four of these questions.

Beginning with the fourth question, the general function of perceptual learning is to improve sensory systems, which is particularly important if there is a malfunction. The processes responsible for perceptual learning do not represent new information from the environment external to the organism, the way world-learning processes do. Acquisition tasks with different goals have likely evolved different mechanisms to achieve them (Rozin & Schull, 1988). Consequently, answers to the other three questions may differ as well.

Continuing backward, learning will manifest itself by changes in what is actually seen, in the case of vision, or heard, felt and so forth in the other modalities. Perceptual learning can be observed if, as a result of experience, the same proximal stimulus leads to a new percept and continues to do so in the absence of new information.[3] In principle, perceptual learning

[2] It is interesting that many examples of "critical periods" for acquisition appear to fall into this third category. It is not clear whether this is coincidental, or whether it reflects something important about learning.

[3] The second part of the definition, "continues to do so in the absence of new information," may not be necessary. Many definitions of learning processes include some mention of stability or relative permanence, and I have succumbed to the same temptation. Definitions are often based on implicitly stated theories (Rescorla & Holland, 1976); the initial division into "experience that makes us better" and "experience that makes us worse" may also capture the intuition of longevity. Things that make us better are usually those that stay around for a while, whereas we hope that things that make us worse are temporary.

distinguishes itself from world learning, which can manifest itself as changes in knowledge, beliefs, and expectations, but which typically will not alter what is perceived. Of course, identifying when perception ends and knowledge begins presents both a theoretical and a methodological challenge. Some phenomena will likely be clear cases of perceptual change within any theory of perception; others will be more difficult.

For instance, if an object looks white under particular lighting conditions before experience, but afterward looks red under the identical physical conditions, then a change in what is seen has taken place. This example will be expanded later in the discussion of the McCollough effect. A phenomenon more in the "gray area" may be habituation. Under some views, the repeating stimulation becomes uninfluential because of a decrease in attention. Whether attention changes should be regarded as changes in perception depends on general theories of perception. A second clear example of a perceptual change occurs if an object is localized straight ahead of the nose before experience, but six inches to the left following experience. The change could occur in vision, so that the object looks as if it is in a different position,[4] or the change could be a different modality. For instance, if the same pattern of proprioceptive joint information produces a different conclusion as to where the arm is after experience, an object will be localized differently via touch. The example of perceived location change following experience will be expanded later in the discussion of adaptation to rearrangements of space.

At this point, however, we encounter a puzzle. The function of perceptual learning is to correct "malfunctions," and its occurrence can be observed by looking for a change in perception. But what could possibly be the circumstances that would cause such a change to occur?

IV. Paradox of Perceptual Learning

Most models of world learning assume that the output of sensory systems is correct, and begin the models from that point. New information about the world is conveyed via the sensory systems. If information about the world comes through the senses, why not always interpret a pattern of stimulation as reflecting the world? How could the mind ever know that the

[4] Such a change in perceived visual location is likely to result from a change in felt direction of gaze, where the same eye position signals (mostly efference information) are interpreted as a different eye position (e.g. Howard, 1982). This in turn affects the seen location of the target, which is dependent on both retinal position and eye position information.

proximal stimulus is due to an internal error, and that it is not simply reflecting another facet of the world?

Part of the answer is that not all knowledge of the world does come through the senses. Roger Shepard argues that enduring regularities in the world are likely to have been internalized by perceptual systems (Shepard, 1984, 1987, 1991, 1992). A few examples are the existence of three dimensions of space, the existence and nature of gravity, the 24-hr light/dark cycle, and the rigid motions of objects. We have evolved to know these things, presumably, because possession of this knowledge conveyed some advantage over having to acquire it during the lifetime of the individual through "trial and possibly fatal error" (Shepard, 1984, p. 432.) Elizabeth Spelke and colleagues (Spelke, 1990; Spelke, Breinlinger, Macomber, & Jacobson, 1992) have elucidated a number of principles, specifically about the behavior of objects, that young infants appear already to know with little or no input. Objects in our world are bounded, travel on continuous paths, do not mysteriously vanish, and so forth.

With knowledge about how the world works, new input encountered naturally or through experimental manipulations can now be in agreement with, neutral in relationship to, or inconsistent with the innate information. The majority of encounters are likely neutral with respect to these constraints, in which case new information is added to the preexisting knowledge base. If the new input is in perfect agreement with the internal knowledge, then there is no need for learning of any sort. If, however, the two sources of information are inconsistent, then there are two possible inferences: Either the constraint is abandoned as false, or the constraint is right and the fault is internal. There are likely a number of constraints never, or almost never, relinquished, such as the three-dimensionality of space. Information inconsistent with these constraints can provide the basis for perceptual learning to occur.

One constraint important for perceptual learning is that an object cannot occupy more than one place at one time (Bedford, 1993). Information discrepant with that a priori knowledge of the world will suggest an internal malfunction. If you localize a desired object a few inches to your right through vision, but simultaneously find the object physically to be somewhere else, then the information from the two modalities taken together has violated the constraint. These are precisely the conditions produced in a prism adaptation experiment, in which the wedge prism displaces the visual locations of objects from their locations as determined by the remaining modalities. The inconsistency will be encountered in nature as well because of physical growth in childhood (Held, 1965). As the shoulders get broader, for instance, the mind's sense of where the body is in space will be wrong. Consequently, that modality will produce a dis-

agreement with vision over the location of objects. The disagreement may also be produced after periods in the dark when vision and proprioception drift with respect to each other possibly because all complex systems drift out of alignment without feedback (Howard, 1982). Two other possible sources of natural discrepancy are the visual distortions produced by different environments, such as lakes and rivers, and by injury (Howard, 1982).

It is typically accepted that the basis for adaptation occurs when there is a "discrepancy" between two cues that determine the same parameter and usually provide the same value (Wallach, 1968). Adaptation is the process by which the discrepancy is reduced or eliminated (e.g., Welch, 1978). The above analysis suggests why a mismatch is problematic. Were it not for the built-in constraint, the visual and motor inputs could be providing new information about the current world. For example, the mismatch could imply that the same object is in two places at once. Instead, the information suggests an internal malfunction. There will be an adjustment to one or both modalities such that the location constraint is upheld.

I turn now to more explicit discussion of adaptation to rearrangements of space. Much of the research described can be viewed as providing answers to the question "What is the content of learning?" for this type of perceptual learning.

V. Prism Adaptation: New Variants on a Classic Paradigm

Previous research on adaptation to rearrangements of one-dimensional space has been largely limited to the nearly uniform displacement produced by the prism. The relevant dimension of space is a set of positions from left to right in a horizontal plane. This physical dimension leads to at least two psychological dimensions: the positions as they are localized visually, and the same positions as they are localized with the hand, or proprioceptively (see Fig. 2). A uniform displacement of all visual positions with respect to motor positions is only one of an infinite number of mappings that can be arranged between the two continua. To describe just a few, there can be a discontinuous mapping where space is split down the middle and only regions to the left of straight ahead are pulled further to the left. There can be a many-to-one mapping, where, in the extreme, every single visual position would correspond to exactly the same location in motor space. This would imply that regardless of where you localized an object visually, you would always reach for exactly the same absolute location to obtain the object. Or we can impose an arbitrary mapping in which visual and motor locations chosen at random are paired! Figure 2

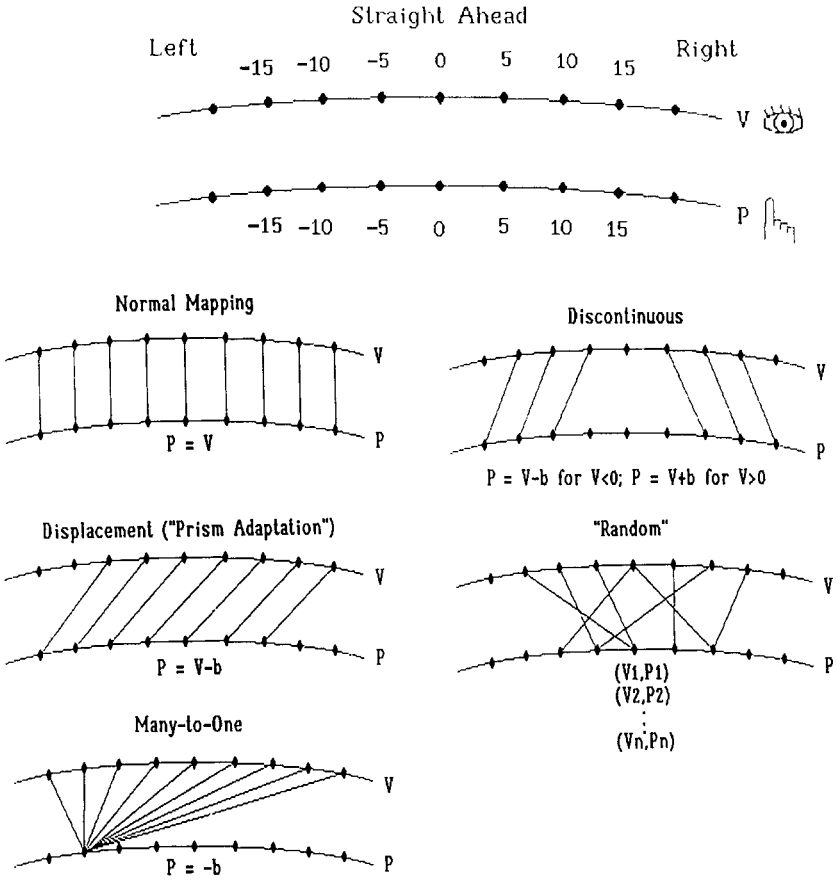

Fig. 2. Some mappings between spatial dimensions. The connecting lines show which positions along the visual dimension (V) go with which positions along the proprioceptive dimension (P) for different mappings. In the equations beneath, "b" stands for a numerical constant.

shows a schematic of the visual and motor dimensions, the normal mapping, the displacement mapping, and the sample mappings.

To fully understand the rules of acquisition, we need to know not just how the system performs with simple lateral displacement but also what it will do for any conceivable mapping that could be imposed, no matter how bizarre or improbable. The research enterprise I have been conducting for the last few years is concerned with the question, What are the rules when two dimensions are to be connected (Bedford, 1989, 1992a, 1992b, 1993)? For instance, should a mapping between two entire dimensions be viewed

as a collection of individual pairs of locations? One of the salient features of coordination is that an individual needs to know where to direct her hand not just for one location in space but for any location where an object may appear. If a mapping can be decomposed into a list of independent visual-motor pairs, then rules of acquisition may be similar to those of simple associative processes. Similarities between adaptation to rearrangements and simple associative learning have been noted (e.g., Epstein, 1976; Taub, 1968; Taylor, 1962); if they prove to be important, our theoretical task becomes simpler.

There are two complementary approaches for uncovering preexisting constraints. One is to provide ambiguous, or incomplete, information about a mapping by specifying the new arrangement for only very few regions of space. The logic of this manipulation is that there are an infinity of mappings that would be consistent with only a handful of visual–motor pairs. The particular interpretation imposed by the perceptual system will presumably reflect the internal rules. An analogy from the domain of natural language may come from the acquisition of word meaning, in which a novel label (e.g., *dog*) could map onto any one of an infinite number of concepts. The interpretation chosen by the child is believed to reflect internal biases, such as the "Whole Object assumption" (Markman, 1990). In the domain of spatial acquisition, we can determine which interpretation was favored by inspecting how behavior generalizes to locations different from those explicitly specified during training.

For instance, imagine you are in a dark room, with a single small target visible in the distance. Through the retinal position of that target, along with eye and head position information, you can recover where the target is located in body-centered coordinates—say straight ahead. Now imagine you reach for the target but discover that if you direct your hand straight ahead, the target is physically not there. Instead, your hand must be a few inches off straight ahead to find the target in motor space. The situation is repeated enough to rule out the possibility of a fluke. For that target, you begin to point to the side of straight ahead, but what should you do if a target appears in a different location?

Describing "one-pair training" in the above fashion is risky because it implies conscious awareness of the mismatch and conscious choice about how to generalize, neither of which need be present. In fact, I will argue later that other contexts, in which conscious intervention is more common, lead to radically different outcomes. This description serves to give a sense, however, of the problem facing the perceptual system if it receives minimal information.

The second approach starts from the other extreme by attempting to provide unambiguous, or complete, information about the nature of a

mapping. This is done by experimentally specifying visual–motor pairs for many points in space, thereby dramatically reducing the set of mappings logically consistent with training. Rather than inspecting generalization, here we look to see how well the novel arrangement was accommodated. Presumably mappings easily accommodated are "natural" in some sense, whereas those that create difficulty go against the internal structure. In addition, the structure may also be revealed by studying the errors that result when a mapping is poorly acquired.

The first strategy is optimal for revealing the hidden biases brought to the task, whereas the second may be more useful for identifying the absolute constraints on learning. Note that the general questions of rules and constraints are different from the research question most frequently asked of prism adaptation: Which system changes? The change in pointing observed following exposure to rearrangements could be due to a change in vision, such that you point in a different location because you see the object in a different location. Or it could be due to a change in body proprioception, so that you point to a different location because your hand feels that it is in a different location. (See Harris, 1965). The study of which type of experience with the prism leads to which outcome has been extensive, and research on this issue continues (e.g., Redding & Wallace, 1988a, 1988b). For instance, moving your hand from side to side while watching through a prism produces a change in the felt position of the hand, whereas seeing the hand only at the end of pointing motion adds a visual change as well. The different possible sites of adaptation to rearrangements and the experiences that lead to them will not be reviewed here; thorough reviews can be found in the works of Harris (1965), Howard (1982), Redding & Wallace (1992), and Welch (1978, 1986).

Before discussing the outcome of manipulations with a few novel pairs (strategy of ambiguous information) and of more complicated mappings defined over all of space (strategy of unambiguous information), I will first briefly describe the methodology. This section has two purposes: to explain the general paradigm and to make clear how the location constraint is violated experimentally.

A. Methodology

As in many prism adaptation experiments and other studies designed to test for effects of experience, the procedure consists of three mains phases: an initial testing phase (known as *pretest* in the prism literature), a training phase (*exposure*), and a final testing phase (*posttest*). The observer sits at the center of a semicircular array of lights, which are located every few degrees (Fig. 3). Lights are illuminated one at a time, and a

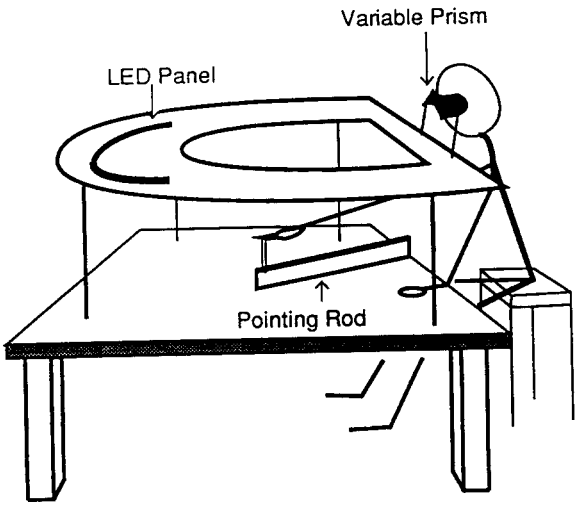

Fig. 3. Apparatus for prism adaptation experiments. From "Constraints on Learning New Mappings between Perceptual Dimensions" by F. L. Bedford, 1989, *Journal of Experimental Psychology: Human Perception and Performance, 15*. Copyright © 1989 by the American Psychological Association. Reprinted by permission.

subject is asked to point to the light. The room is dark, and pointing must occur without the benefit of visual guidance. The purpose of this phase is to provide a measure of how accurately subjects point before receiving any training.

The training phase is especially designed to allow the two strategies to be implemented. The subject wears a single light on the finger, and, as before, a light in space is illuminated in the darkened room. The subject is told that if she points accurately to the light in space, the light on her finger (*finger LED*) will light up. Varying arm position by a fraction of a degree on either side turns the finger LED off again. The subject's task is to light the finger LED and keep it lit as much as possible during each trial.

If vision were undistorted, the task would be especially simple. Even without visual guidance, it is easy to direct the arm to a target in space, an action that would cause the finger LED to turn on. But the subject looks through a prism, making the target appear to the side. Therefore, the subject initially points to the side of the target's true location, failing to illuminate the finger LED. She is encouraged to explore, and usually with one or two attempts finds the correct position, which causes the finger LED to light up. (See Fig. 4, top two panels; the third panel is discussed later in the article.) The subject is misinformed that it is "disorienting being

PERCEPTUAL LEARNING

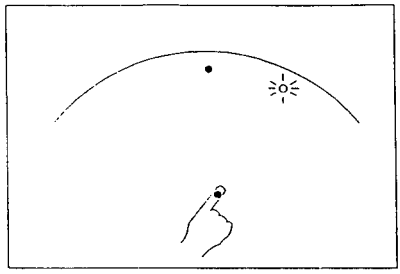

Initial Attempt

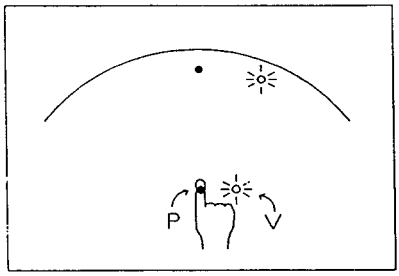

Successful Attempt

COGNITIVE LEARNING

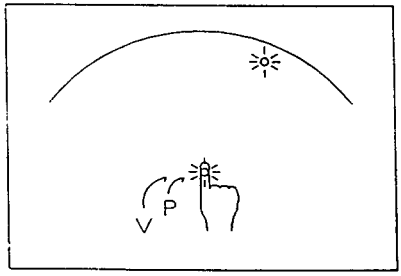

Successful Attempt

Fig. 4. Pointing with the finger LED. Top two panels show an initial and a successful attempt, both made while looking through a prism (Perceptual Learning). The bottom panel shows a successful attempt using a different "Cognitive Learning" paradigm. V, visual position of the finger; P, felt (proprioceptive) position of the finger. From "Perceptual and Cognitive Spatial Learning" by F. L. Bedford, in press, *Journal of Experimental Psychology: Human Perception and Performance*. Copyright © by the American Psychological Association. Adapted by permission.

in the dark," a small deception that she readily believes to be the source of her difficulty. Further practice is usually required to keep the finger LED on steadily, a difficulty that keeps the subject's attention on the task. When the finger LED is successfully illuminated, the subject is localizing the light on her finger through proprioception in one location, but sees the light (also through the prism) in a different location. This situation creates the apparent violation of the constraint that an object can be in only one place at one time, because the finger is detected in two different locations.

The paradigm essentially isolates individual "points" in space because the room is completely dark and visual feedback is provided for only a very small region of the finger. A successful illumination of the finger LED provides information about only one visual–proprioceptive pairing. Environmental input can therefore be limited to only one, or a few, locations, enabling the research strategy of ambiguous information to be easily implemented. This arrangement differs from the typical training procedure used in adaptation experiments in which the mapping is specified over a large continuous range of locations, for example, watching one's hand while pointing to targets in full-room illumination. In addition, mappings other than uniform displacements could be created by using a variable prism under computer control. The prism enabled lateral shifts of anywhere from 13° to the left to 13° to the right. Choosing different shifts for different locations permitted complicated mappings to be created.

Following training, testing occurs again. The change in pointing from before the experience to after the experience serves as the measure of acquisition.

B. Ambiguous Information

What happens when the input is limited to only one or two points? I find that following repeated training trials at only one location, behavior generalizes in a rigid fashion (Bedford, 1989). That is, however far you pointed off from the trained location, you point that amount off for all positions, even those a fair distance away. This appears to be true regardless of where the trained location occurs, although central training produces larger overall shifts. Figure 5 shows the change in pointing for two training conditions, one when the single location was straight ahead, and the other where it was not. The x axis shows the position of the test target, from left (−) of straight ahead (0) to the right (+). The y axis shows the change in pointing from pretest to posttest. If the experience had no effect, there would be a change of 0, and the data would lie along the 0-degree horizontal line. Input at one location proved sufficient to affect both that location and other locations.

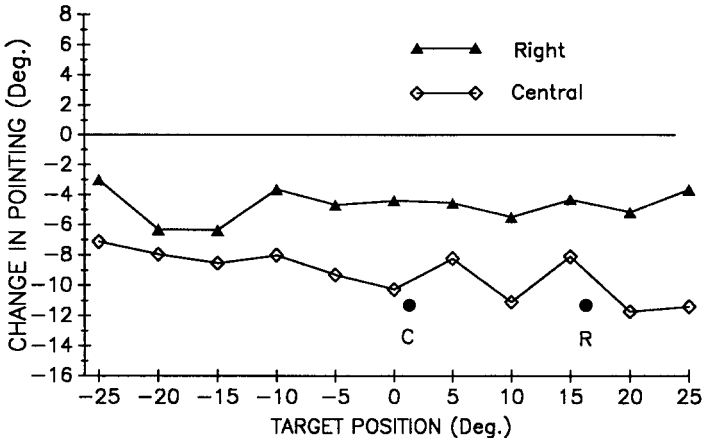

Fig. 5. Training with one pair located either straight ahead (C) or to the right (R). Location of training pair shown by filled circle. From "Constraints on Learning New Mappings between Perceptual Dimensions" by F. L. Bedford, 1989, *Journal of Experimental Psychology: Human Perception and Performance, 15*. Copyright © 1989 by the American Psychological Association. Adapted by permission.

What if training is inconsistent with a rigid shift for all of space? Consider the additional information provided by a second target in a new location. The second pair can be chosen so as to be consistent with a rigid shift, or to violate that solution. If touch is offset from vision by exactly the same amount for both locations, (e.g., 10° to the right) then interpolation to intervening locations is rigid. The result is not unexpected, given that only a single pair is needed to produce a uniform shift. If the offsets are in opposite directions (e.g., when the target is to the left, point further to the left, but when it's to the right, point further to the right), then behavior can be described as a rescaling of space. One dimension gets uniformly stretched with respect to the other. The intervening locations are filled in accordingly, even though the mapping was never encountered at those locations (Fig. 6).

The consequences of one- and two-pair input argue for two main conclusions. First, a mapping between visual space and motor space cannot be thought of as a collection of individual pairs of locations. If it were, the pairs would be largely independent and we would expect the system to generalize differently. We would be more likely to obtain the familiar generalization gradient surrounding each trained pair. Each trained location would provide minimal information about others, and the impact would decline rapidly with increasing distance from the trained location.

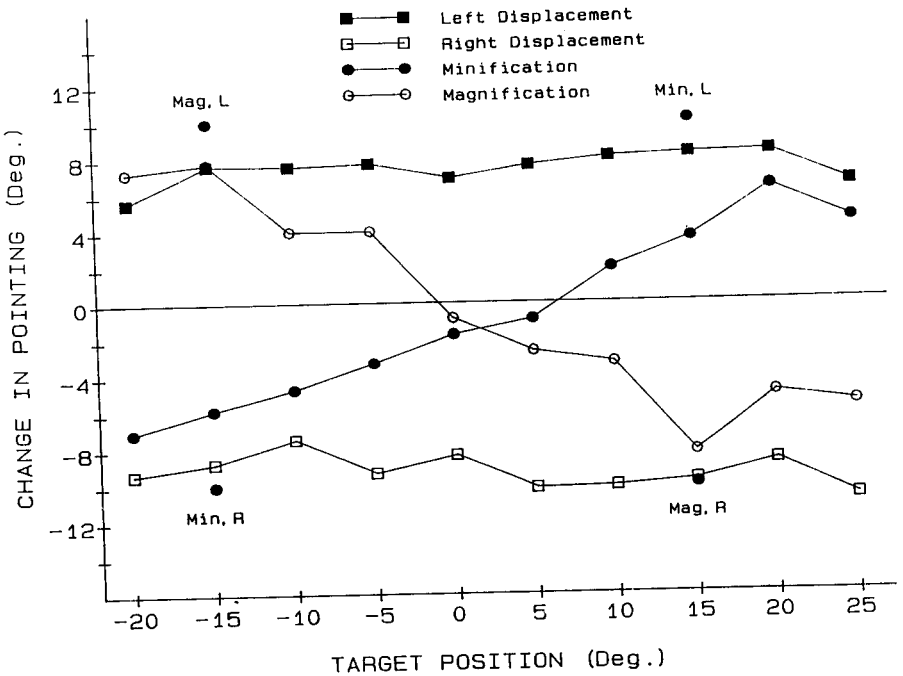

Fig. 6. Training with two pairs—interpolation (Bedford, 1989). Left Displacement (L): Both visual positions shifted to the left. Right Displacement (R): Both visual positions shifted to the right. Minification (Min): Left visual position shifted to right; right visual position shifted to left. Magnification (Mag): Left visual position shifted to left; right visual position shifted to right. From "Constraints on Learning New Mappings between Perceptual Dimensions" by F. L. Bedford, 1989, *Journal of Experimental Psychology: Human Perception and Performance, 15*. Copyright © 1989 by the American Psychological Association. Reprinted by permission.

In the intuitive description given earlier for one-pair training, this would be the equivalent of adopting a strategy, perhaps wisely, of not generalizing the learned response to locations far removed from the conditions of training. Another way to think of the failure to get independence is that training two pairs simultaneously would not always be predictable from training each pair by itself and combining them additively.

Instead, training seems to influence the entire dimension. Note that information at only a single location was needed to drive the change for all of space. Mechanisms for which only minimal input is needed from the environment for acquisition are usually indicative of highly constrained, rule-governed, specialized learning organs.

Second, the "constraint" revealed by these manipulations can be described as *linear*. Of the infinite number of possible mappings consistent with the limited input, linear mappings were chosen. The generalized pattern always conformed either to a rigid shift (intercept parameter) or to a uniform rescaling (slope parameter). Note that within linearity, intercept changes seem to be preferred. Scale changes occurred only with information inconsistent with a change in a single intercept parameter.

C. Extrapolation

We were struck by an unexpected feature of the data. Although interpolation between trained positions conformed to a sloping line, when an intercept shift was not possible, the same was not obviously true for positions that did not fall within the two end points of training. Rather than extrapolating along the fitted line, pointing seemed to level off, perhaps even to drop back (refer to Fig. 6). Figure 7 shows the outcome of some manipulations tried in order to further explore the issue.

1. Absolute Size

Is the result simply a performance limitation? Perhaps the maximum permissible shift in pointing is something like 9°. If so, requiring anything greater than that amount, as might be required by extrapolation, would not be seen in behavior. This possibility proved insufficient. A two-pair training situation was designed such that the extrapolation, if it occurred, would fall within the 9 or so degrees known to be permissible based on prior experiments. Figure 7A shows the training pairs and the data. In addition, the dotted line shows the "interpolation line," calculated from the data of the three test locations that are within the two training end points. Here too, extrapolation along the same line that is interpolated does not occur, suggesting that this is a general phenomenon.

2. Centrality

We investigated the conditions under which extrapolation may be more likely to occur. Thus far, extrapolation would require generalizing from two points straddling a central region (e.g., 7° to the left of straight ahead and 7° to the right) to a non-central region (e.g., 7° to 25°). Yet many systems are known to be linear only over a central range. Perhaps this knowledge has been internalized in some sense, so that a linear function in the middle of a spatial continuum would not be generalized to a non-central region in the absence of solid proof. It may be easier to go in the other direction.

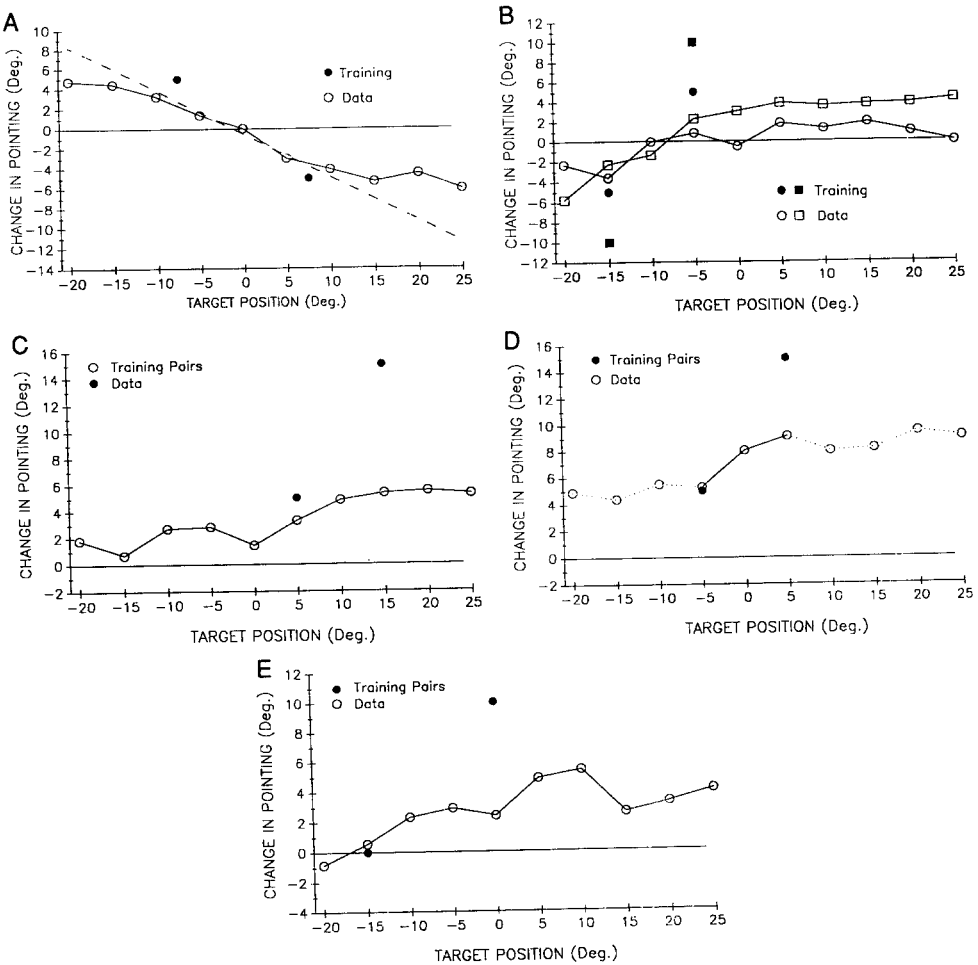

Fig. 7. Training with two pairs—extrapolation. Dashed line in panel A shows interpolation line fit to the three test locations between the two training positions. A: From "Constraints on Learning New Mappings between Perceptual Dimensions" by F. L. Bedford, 1989, *Journal of Experimental Psychology: Human Perception and Performance, 15*. Copyright © 1989 by the American Psychological Association. Adapted by permission. E: From "Perceptual and Cognitive Spatial Learning" by F. L. Bedford, in press, *Journal of Experimental Psychology: Human Perception and Performance*. Copyright © by the American Psychological Association. Reprinted by permission.

The two training pairs were "moved" to the side, such that both visual locations were to the left of straight ahead. Figure 7B shows the pattern of generalization. The primary outcome is that extrapolation along the interpolated line does not occur to central regions either. This result occurred

for two different sets of training pairs, one with 5° offsets and the other with 10° offsets.

3. *Relative Size*

Intuitively, the most difficult requirement of extrapolation may be having to infer a change that is *larger* than any experienced during training. If the two trained locations required pointing shifts of 5 degrees (one leftward, and one rightward), it may be incautious to assume that some locations must require pointing shifts of 6° or 8°, or even more. One difficulty is that there is no natural stopping point: At what value would the extrapolation end? What if the two input pairs required a different best-fit line, which, if extrapolated, would produce changes in pointing that were only smaller than those encountered in training?

Instead of moving the two pairs to the side, they were slid upward. That is, instead of the two offsets between visual and proprioceptive positions going in opposite directions, the offsets of the pairs were in the same direction but were of different sizes. In one of the manipulations, the two visual locations were on the same side of straight ahead, to assess the combination of this factor along with "centrality" (Fig. 7C). In another, they were on opposite sides of straight ahead. The data are most striking in Fig. 7D, where extrapolation to the right would require a larger response and to the left a smaller response. There appears to be little difference: Extrapolation in a downward direction does not occur either.

4. *Normal Experience*

Finally, a condition was chosen that seemed certain to produce a different pattern of generalization. If you receive training at two pairs, so that the offsets are in opposite directions, or so that the offsets are in the same direction but have different amounts, it is arguably difficult to know how to go beyond the specific training instances. However, suppose one of the two pairs had no offset at all. For instance, when a target is localized straight ahead, you must reach slightly to the right, but you also observe that when it appears somewhere else (e.g., 15° to the left), you just reach directly at the target, just as you always do. How do you extrapolate?

It seems that the most conservative strategy is to assume that, unless there is information to the contrary, the straight-ahead position is simply the "odd man out." Whereas the other manipulations were conditions chosen as most likely to produce linear extrapolation along the fitted line, this manipulation was chosen in order to have the opposite influence, that is, to produce more standard generalization decrements.

This did not occur (Fig. 7E); if anything, there was more extrapolation here than in any of the other conditions!

Overall, the pattern of data suggests that a sloping line fit between two points is not extended to locations outside that range, even under conditions in which it seems sensible and/or easy to do so. Nor does pointing drop back toward zero, even when that decision would be conservative. Under a variety of different conditions, the data may best be described as "leveling off," so that however much the system shifted at an end point, that same shift is continued for all locations beyond. For some of the conditions, there appears to be a little extrapolation along a line with a shallower slope than that of the interpolated line. Part of this effect may be due to consequences of presenting visual stimuli all to one side of straight ahead, or of presenting the two training pairs very close together. Whether some minimal extrapolation remains has not yet been determined. Extrapolation experiments are difficult to do. They are more likely than interpolation experiments to produce a few subjects with wildly aberrant means and variability. At present, the data appear to reflect a tendency for the system to extrapolate with caution.

D. Unambiguous Information

I turn now to a few manipulations using the opposite strategy. To present "complete" mappings, typically approximately nine distinct locations were chosen. The experiments reported do not reflect a systematic exploration, but rather loosely follow one particular line of thought.

1. Isolation

The type of generalization following degraded input suggests a preference for modifying all the locations along a spatial continuum, rather than only one. As discussed, neither the presentation of a single shifted location nor the special case in which a second location was consistent with the normal mapping led to the modification of only one small region of space. However, these results do not prove that a dimension *can't* be broken into small independent regions. If breaking a dimension into small independent pieces is less preferred than are linear functions, that ability could easily be hidden under circumstances in which the decision is left up to the system. To determine whether one location could be singled out, a mapping was created for which just about any interpretation other than "only one region has changed" would be a failure to learn correctly.

When a target appeared straight ahead, its physical location was actually to the right, but when it appeared anywhere else, coordination could proceed normally. Figure 8 (top) shows the outcome following training

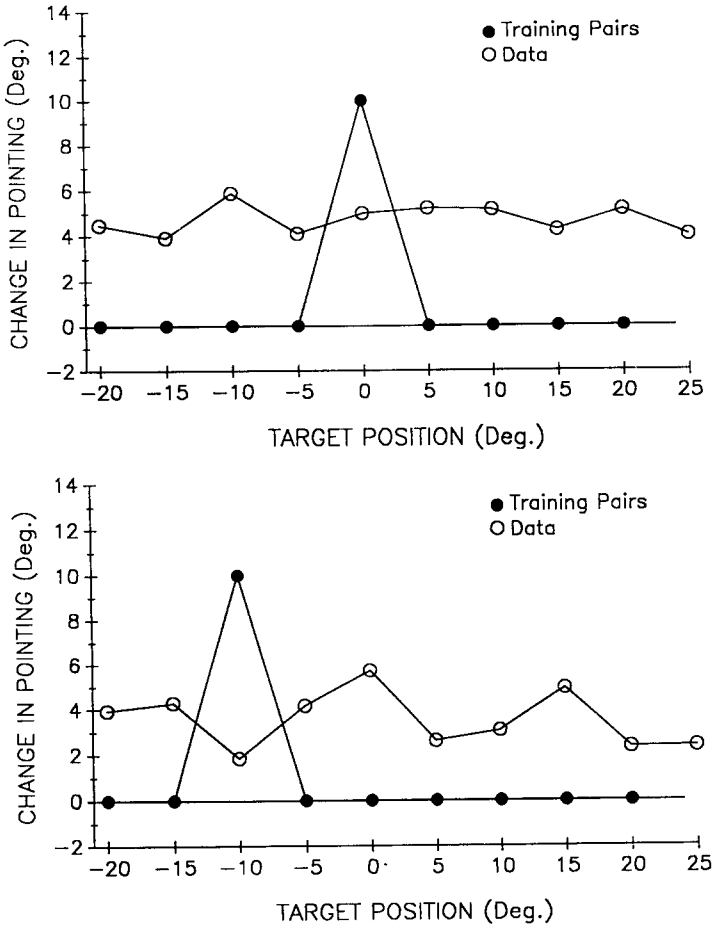

Fig. 8. Training with many pairs to isolate one region. Atypical location either straight ahead (top panel) or to the left (bottom panel). Top: From "Perceptual and Cognitive Spatial Learning" by F. L. Bedford, in press, *Journal of Experimental Psychology: Human Perception and Performance*. Copyright © by the American Psychological Association. Reprinted by permission.

with the one location shifted, and the remaining locations unchanged. The system was, in fact, unable to accommodate this intuitively simple transformation. If a mapping was really a collection of independent pairs, this particular transformation would be easily learned. However, it appears unlikely that the dimension can ever be divided, even under conditions in which it is highly adaptive to do so.

That which subjects impose on this mapping converges with what they interpolate following degraded input. Though the mapping was not learned correctly, pointing did change. There was a roughly uniform shift for all locations, by an amount roughly in between the 10° offset and zero. This is precisely the best-fit linear function under the circumstances (50% of the training occurred at the straight ahead position, and 50% was equally distributed among the remaining locations). Unfortunately, the fit is not always perfect. Figure 8 (bottom) shows what happens following similar training, except where the odd-man-out target is off to the side rather than at straight ahead. Whereas the best-fit linear function would be a sloping line, the change was instead a rigid shift. This may be another manifestation of the greater preference for intercept over slope changes, as seen with one- and two-pair training. Clearly this issue needs further investigation.

Why wasn't the intuitively simple isolation mapping acquired? One possibility is that only linear functions can be used to realign the two spaces. A second possibility is that the mapping was particularly difficult because it violated topological properties, as well as linear ones. Shepard (1989) while referring to neural representations and models, suggests that a rearrangement of space that disrupts topology is not likely to be learned because that would ". . . completely defeat the local connectivities of the topographically organized system" (p. 128). Figure 9A shows how locations in visual space get mapped into locations in motor space for the mapping where only the straight-ahead position was changed. Note that the property of being a one-to-one mapping has been destroyed, as has the initial order of points, both of which are topological properties.

2. *Topology*

Are topological properties relevant for realigning dimensions? Two mappings were devised that used identical visual locations and identical offsets, except that one of them violated topology and the other did not. For one group, all visual locations to the left of zero required pointing further to the left by 10° and all locations to the right required pointing further to the right by 10°. The straight-ahead position remained normal, that is, a 0° offset. For the other group, the directions were reversed so that all visual locations to the left of 0 required pointing further to right and locations to the right required pointing further to the left. As in the other group, straight ahead was unchanged. Figure 9B shows the distortion of space produced by these mappings; it can be seen that the uncrossed group is a one-to-one mapping with original order preserved and that the second, crossed, group is not. For instance, in the crossed group, if you wanted to obtain physically a stimulus that appeared either straight ahead, 10° to the left, or 10° to

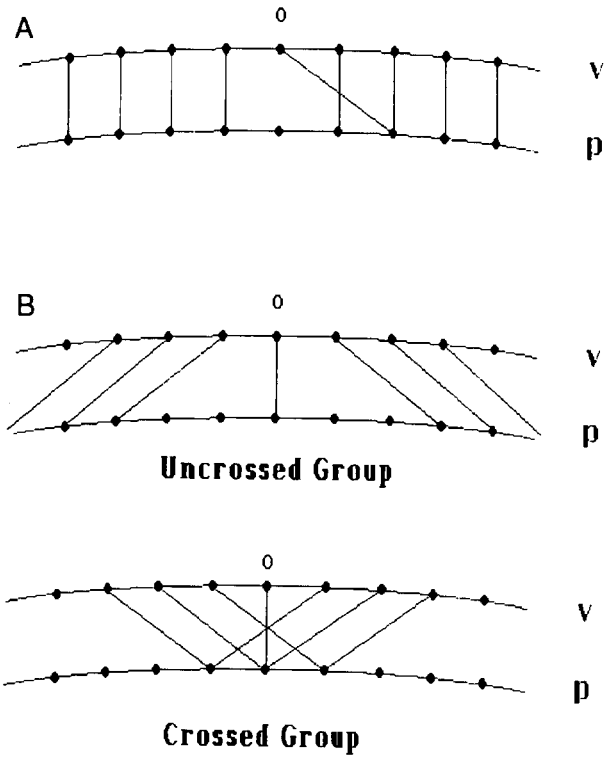

Fig. 9. Mapping required in the isolation (A) and topological (B) experiments. v, Positions along the visual dimension; p, positions along the proprioceptive dimension.

the right, you would do so by reaching straight ahead for all three locations. In the uncrossed group, the distinctness of these three locations remains.

The data are presented in Fig. 10, which shows how pointing changes as a function of target position in the two groups; the figure also shows a comparison of the size of slope changes. The changes in slope are different in the two groups, with the uncrossed mapping (.28) producing more than double the amount of change than does the crossed mapping (.11). (Both slopes are significantly different from 0, and from each other.) Slopes were compared because, interestingly, only linear components reached significance when the changes in pointing were tested for polynomial components up to Degree 9. The mappings would be more accurately captured by a cubic, which would reflect two "bends" in the data. The data for at least the uncrossed group look as if they were heading in that direction.

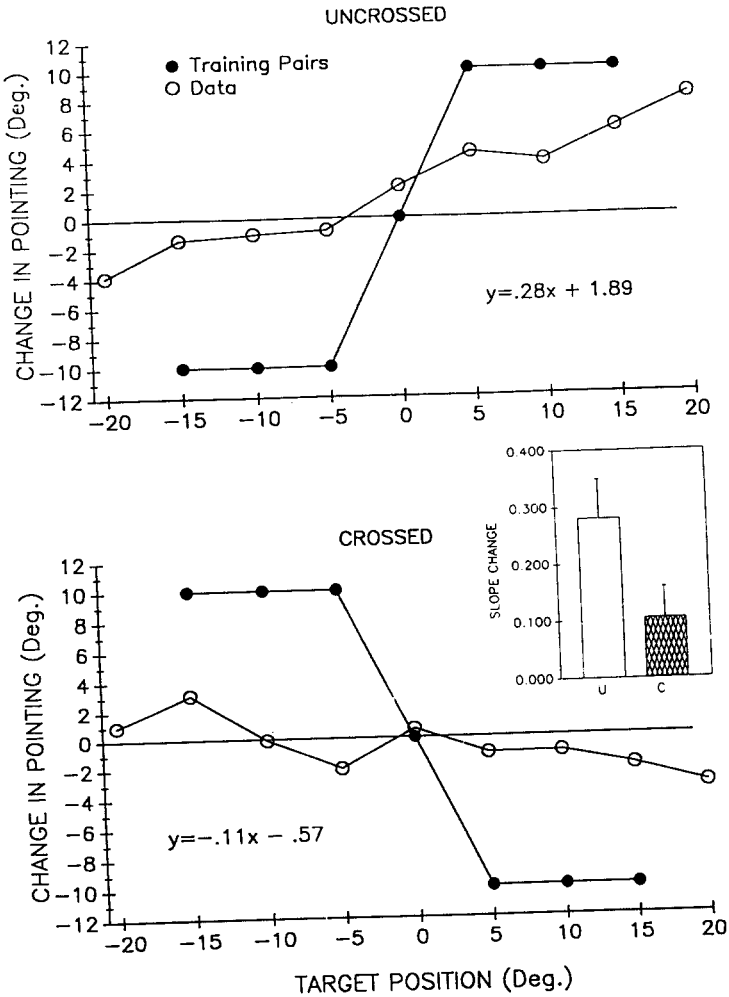

Fig. 10. Data for topological experiment. U, Uncrossed group; C, crossed group.

Whether more training with the uncrossed mapping would produce more accurate results remains to be determined. It is possible that the preferred linear function is present in intermediate stages and then gets refined accordingly. Because statistically there is evidence only for a linear fit, whether transformations that are nonlinear but that preserve topological properties can be accommodated is not yet known. However, either way,

there clearly is a sensitivity to topological properties, as demonstrated by greater change to mappings that are one to one and that preserve order.

Does the fact that there was some change in the crossed group suggest that even gross distortions of space can be accommodated? The answer is no. Figure 11 shows the mappings actually acquired, in which it can be seen that for the crossed group as well as the uncrossed group visual–motor behavior never became many-to-one, nor did it switch order. There appears to be a sensitivity to topological properties, so that mappings that do not violate topology are more easily dealt with. In addition, there is so far no evidence that violations of topology are ever acquired. Whether some topological properties are more influential than others is currently under investigation.

3. *Discontinuity*

A different strategy for dealing with the uncrossed mapping would be one in which the system split up the space, so that the left half of space were

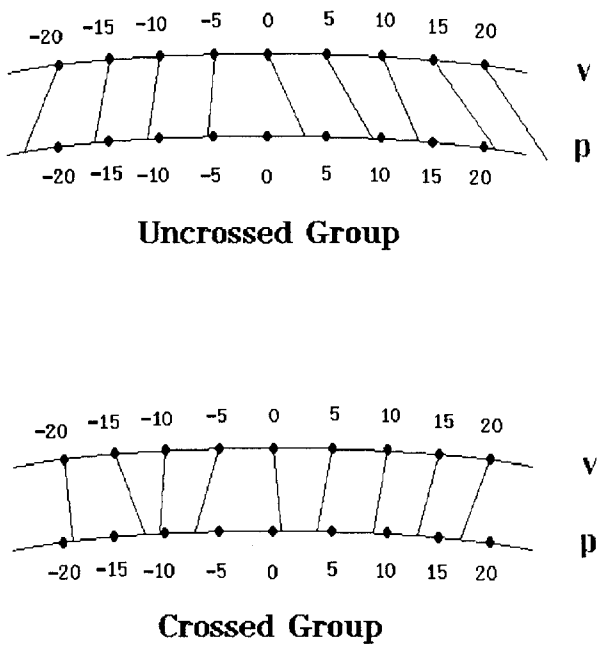

Fig. 11. Mappings acquired in the topological experiment. v, Positions along the visual dimensions; p, positions along the proprioceptive dimension.

pulled further left, and the right half further right. This would be a nontopological transformation because destroying the continuity of points, like destroying distinctness and order, destroys the topology. However, the system is good at rigid shifts, and in addition, more than one adaptive state can be present at the same time when they are contingent on different cues (e.g., Hay & Pick, 1966; Kohler, 1964; Taylor, 1962; Welch, 1971; see Welch, 1978, pp. 96–99 for a review.) It might be possible to code the uncrossed mapping as two intercept parameters, each relevant for a different half of space. This solution may have been discouraged by the visual–proprioceptive pair for which the straight-ahead position remained normal (0,0).

Consequently, the mappings were repeated, except with the (0,0) pair omitted. Statistically, the results were little different from the prior experiment. A linear, but not cubic, component was present for the uncrossed and for the crossed group. Like the other study, the uncrossed group showed greater change than the crossed group (slopes = .24 and −.14, respectively) reflecting again the sensitivity to topological properties. Although it is difficult to rule out definitely that separate intercept parameters were fit to the two halves of space, the data do not overwhelmingly support this solution. It may well be that some type of external cue is needed upon which different adaptive states must be contingent, and that space is not easily split otherwise.

E. A Dissociation

What makes this learning perceptual? As discussed earlier, location mismatches between vision and touch for the same object suggest an internal malfunction because internal knowledge about objects has been violated. The inputs that make up the visual-motor mappings are used along with preexisting structure to correct the internal malfunction as best as possible. The research program described suggests some rules governing the acquisition of the new internal states. We can now ask whether the rules are different from learning processes that function to apprehend new information about the world. Another way to phrase the question is whether there is a dissociation between perceptual and world learning on the issue of "the content of what is learned" in addition to the issue of the circumstances that produce learning.

For two of the permutations explored and described above, we directly compared the outcome to a "cognitive" learning task with the same formal properties (Bedford, in press). Would people be able to learn the mapping where only one region had been changed, when not under the control of

the perceptual learning system? In addition, how would they choose to generalize to novel locations following two-pair training?

We needed a task that presented the same mappings but that removed the internal inconsistency between sensory systems. Instead of the prism, the visual–motor offsets were produced through computer software. For instance, a target light located straight ahead would require the arm to be 10 degrees to the right to illuminate the finger LED. Contrast this with the perceptual learning procedure, which also requires that a target appearing straight ahead be pointed to 10° to the right. In that procedure, when the subject is successful, the "hand" looks as if it's also straight ahead, and looks as if it's pointing to the target. This situation occurs because the target's actual position is at 10°, and both it and the hand are seen through the prism. In the cognitive procedure, the hand looks simply where it actually is, at 10°, and it looks as if it is pointing to the side of the target. Because there is no prism, there is no discrepancy between where the finger feels it is and where it appears to be. (Compare the second and third panels of Fig. 3.) Yet, the visual–motor pairings can be made to match any permutation created with the perceptual learning procedure. For the cognitive paradigm, subjects were told from the beginning that this was a learning task. A subject was told that she had to try to figure out and to remember where to put her arm for each different light, and that she would be tested on it afterward.

Without the need for internal correction, and without the ensuing motivation for perceptual learning, the consequences of the particular mappings were quite different. People were able to learn the mapping when one region was shifted and eight regions required accurate pointing (Fig. 12, top), consistent with intuition that this should be an easy rule to acquire. Even more notable was the different pattern of generalization when one region was shifted and one region required accurate pointing. This "discrimination" experiment produced a generalization gradient of the sort more familiar to learning researchers, in which a newly acquired response is used less and less as a function of distance from the explicitly trained stimuli (Fig. 12, bottom).

In these tasks, some type of "cognitive" abilities were used to acquire new information about the environment and to go beyond the exact training stimuli. These processes appear to allow spatial continua to be broken into individual pieces. Models assuming that associations between individual stimuli are the correct units of acquisition may be applicable under these circumstances. This is in contrast to the perceptual learning system, which seems constrained to operate on the dimension as a whole.

One clarification may be useful. In prism adaptation and its variants, the subject's behavior, *as viewed by an outside observer,* is to point to the side

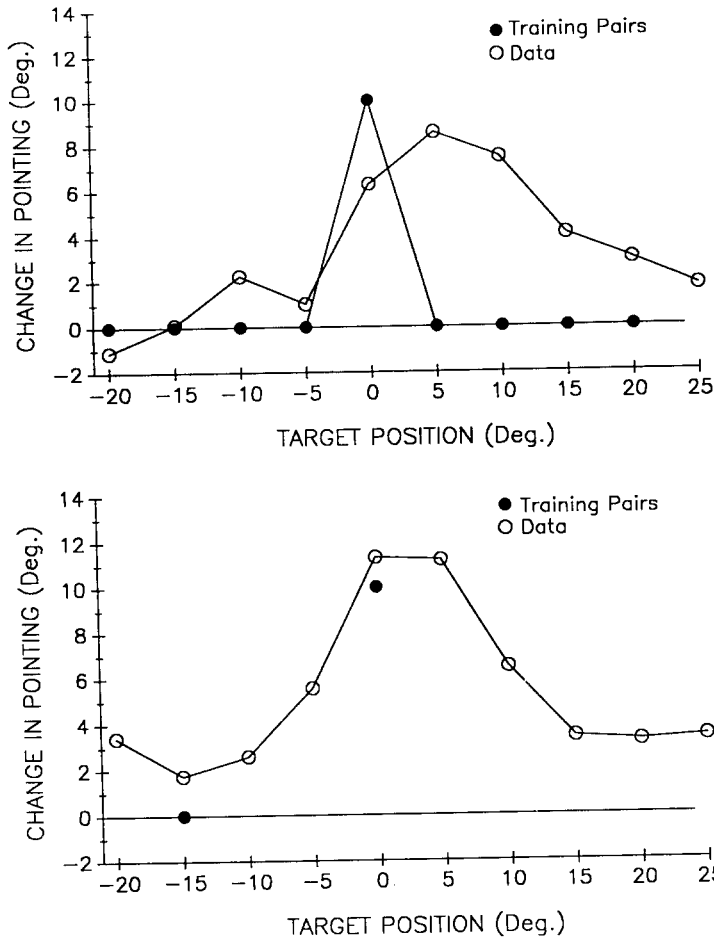

Fig. 12. Data from cognitive learning paradigm. Top panel shows many-pair training. Bottom panel shows two-pair training. From "Perceptual and Cognitive Spatial Learning" by F. L. Bedford, in press, *Journal of Experimental Psychology: Human Perception and Performance*. Copyright © by the American Psychological Association. Reprinted by permission.

of a target's true location. This behavior appears identical to what a subject would do in the cognitive learning task discussed above. However, in perceptual learning there is never a representation that objects are to be found in different places with the visual system and with the motor system. The constraint that an object can be in only one place is upheld. Instead, there is a change within one or both systems, so that the two sensory

systems again coincide. A subject looks as if she points to the side because her hand's position is mislocalized as further to the side or because the target appears to the side. As far as the perceptual system is concerned, she is pointing directly at the visual location. This is in contrast to the cognitive learning context, in which presumably there is a representation to point to the side of the targets seen position.

Another potential difference between the two types of learning concerns the role of conscious awareness. Subjects need not be aware of any discrepancy in order to adapt to rearrangements of space (see Welch, 1978), whereas the cognitive task that we used forced subjects to be aware that there was something new to be acquired. How awareness influences perceptual learning is currently under investigation.

F. Summary of Rules

How can we characterize the rules by which the perceptual system accommodates new mappings? Using the strategies of ambiguous and unambiguous information, we have thus far uncovered the following biases:

1. Acquiring a mapping is not a collection of individual visual–motor associations, but instead involves connecting entire dimensions.
2. There is a preference for linear functions to relate the two continua.
3. Within linearity, there is a preference for intercept changes over slope changes.
4. Extrapolation of a fitted function occurs cautiously, if at all.
5. There is a sensitivity to topological properties of space.
6. There is no evidence that space can be broken into "left half" and "right half."

Linearity clearly plays an important role. We are pursuing the possibility that rigid shifts and uniform stretches and squashes are precisely the internal changs that would be needed to correct for growth of the body in childhood, or the drift between sensory systems in adulthood.

VI. Hierarchy of Transformation Geometry

While linearity is important, it may not be sufficient. For instance, it does not account for the sensitivity to topological properties, nor can it by itself explain the preference for intercept over slope changes. One possible, admittedly speculative, integration of the findings comes from transformation geometry. In brief, the view is that a mapping should be considered

more difficult than another if more geometric properties are altered by the mapping.

The transformation approach to geometry originated with mathematician Felix Klein who showed that different geometries could be ordered on the basis of the number of properties remaining unchanged as a result of a group of transformations (e.g., 1893/1957). The familiar Euclidian geometry concerns all those properties left intact by transformations that move a form in its entirety. A square that has been slid over changes nothing about the "squareness," but alters only its location. All the properties that remain intact—angle, parallelism, size, distance between two points, and so forth—are properties of Euclidian geometry (laboriously studied in high school). Those altered—absolute location—are not in the study of that geometry. Thus, a square, and a square 3 in. away, are equivalent forms in Euclidian geometry. But there are other, more radical, transformations that can be applied that lead to other geometries. For instance, consider pulling on one side of a square to turn it into a rectangle. Such a transformation alters more than location—for example, it alters size and distance between points. Only those properties that remain intact are part of this geometry, for example, parallelism and the property of being a straight line. Whereas in Euclidian geometry a square and a rectangle are distinct individual forms, in affine geometry they are equivalent. Klein laid out a series of transformations, each more radical than the last, and showed how each was the basis of a different-sized geometry. Although his concern with the transformations was the resulting geometry, it is the transformations themselves that most concern us here.

We can imagine visual space as a form, and different mappings as transformations on the form, which turn it into motor space. The more geometric properties altered by the transformation, the more complex the mapping. In the two-dimensional hierarchy, as devised by Klein, there are five distinct levels. They are shown in Fig. 13. What follows is an informal description of the different levels of transformations. Good formal discussions can be found in Modenov and Parkhomenko (1965). The simplest possible transformation picks up the visual space in its entirety and puts it elsewhere, altering only the absolute location. For instance, a square could be rotated or displaced to another location. This as described above as the group of transformations underlying Euclidian geometry. The next, slightly more complex, "similarity" transformations allow uniform expansions and contractions to both dimensions equally. That is, they alter the property of size as well as location, but all the remaining shape properties are unaltered. A square could be transformed into a larger square, or a rectangle into a larger rectangle with unchanged width-to-height ratio, and so forth. Next are the Affine transformations, also mentioned above. Here

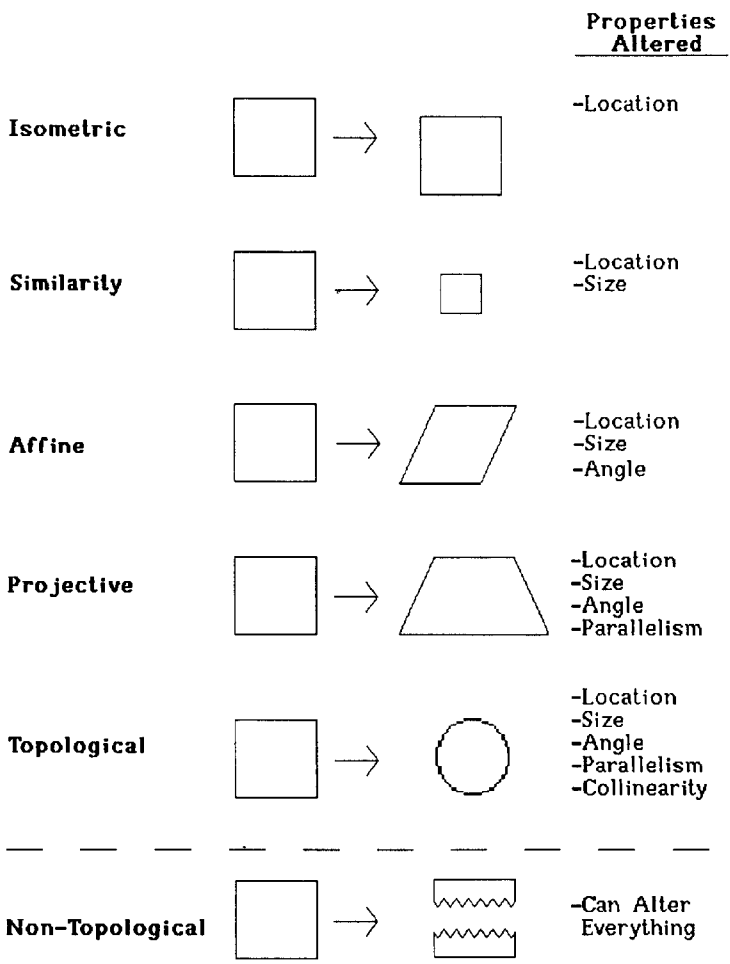

Fig. 13. Sample transformations within Klein's (1893/1957) hierarchy of transformations for two-dimensional space.

uniform expansions and contractions can be applied separately to two orthogonal dimensions. These transformations can turn squares into rectangles, or into rhombuses. Typically, they alter the property of angle, while leaving other shape properties intact, such as collinearity (straightness), parallelism, and order. Affine transformations are the most general type of linear changes. The fourth group are projective transformations, which alter the parallelism of sides as well. Squares, for instance, can turn into trapezoids. Finally, we have topological transformations. These are

the most radical because all transformations of the original form are allowed, provided that distinct points are not glued together (1-to-1) and that continuous points are not split (continuity). These transformations have the power to turn squares into circles, preserving little of what we think of as form, except aspects such as the property of being a closed curve.

For one-dimensional space only, the second, third, and fourth levels collapse to a single level, which we will refer to as "affine." Clearly, a proper test of this view must come from visual–motor mappings in two dimensions of space. This consideration in part motivated the paradigm to be discussed next. Yet from the data that we do have in one dimension of space, the outcome appears consistent with the different levels. The isometric level corresponds to rigid shifts, which we found to be the most highly preferred. Affine transformations were next preferred, which in one-dimensional space correspond to slope changes. These were easily accommodated when the input was inconsistent with a rigid shift. Uniform stretching was, in turn, easier than any topological transformation, reflected by the tendency to fit straight lines to nonlinear mappings. Next, we saw that mappings preserving topological properties produced more change than those that did not. Finally, there was no evidence that nontopological properties can be acquired.

Why pursue this hierarchy? First, if nothing else, it is useful to have a taxonomy of some sort to guide research. There are an infinite number of possible mappings, and clearly no research program could try them all. The existence of a manageable number of distinct classes of mappings allows systematic exploration.

Somewhat more strongly, it should be noted that these different levels are not arbitrary classifications. The different levels are different *groups*, in a mathematical sense. For instance, if one transformation produces a particular mapping, in order to be a group there must also be a transformation that restores the original (presence of an inverse). Although going through the requirements to be a group is beyond the scope of the chapter, the major point is that the concept of a group is mathematics' way of carving nature at its joints. It may be the mind's way as well.

While the different levels must each be a group, that still leaves an infinite number of ways of parsing mappings. Why this one? One based on geometry seems a better candidate than do other schemes, such as polynomials of increasing order. Geometry is the study of form. To the extent that space can be considered to have form, formal properties of geometry are relevant to space and therefore potentially to the perception of space.

There are, however, an infinite number of geometries and corresponding transformations in addition to those described above. Mathematical schemes are "restricted by no other rule than that of avoiding contra-

dictions" (Cassirer, 1944, p. 4). What makes Klein's levels special? While the geometries described are not all Euclidian, they are all within the Euclidian framework. That is, they all share certain fundamental axioms, none of which violate the way our world works (Euclidian with the most axioms, and topology, the least). The correct reflection of our world is not true of other goemetries, for example, that of Reimann. Cheng and Gallistel (1984) argue that although an organism may not represent all properties of the world, it would be crazy to have evolved a rule that gets things systematically wrong. Thus, while mathematics is not restricted by the way the world works, reality greatly limits which schemes are likely relevant for psychological problems.

One final issue on which I will comment concerns the direction of preferences (assuming that the levels are in some sense internalized). For instance, why assume that all of space is a rigid shift (isometric) unless given information to the contrary? Why not assume some other level? A useful analogy here may be the subset principle (Berwick, 1985) in language acquisition. If you have two languages, one of which is the subset of another, children assume the smaller language first and switch to the larger grammar only when given information inconsistent with the smaller one. The logic behind the principle is that it is easier to switch from the smaller to the larger universe than it would be to do the opposite, given the kinds of inputs for language. Children typcially get samples of correct language (*positive evidence*), but don't hear samples of ungrammatical utterances (*negative evidence*). To switch from the subset to the superset, you need only hear a grammatical sentence inconsistent with the subset. But there is no obvious way to get from the superset to the subset, because nothing tells you that you are wrong.

In the current domain, each level of transformations is a subset of all the ones above it. For example, all affine transformations are also topological transformations, but not all topological transformations are affine. While the subset rule is usually thought of as specific to the specialized language-acquisition module, the two domains face similar problems that may have led to similar solutions. As in language acquisition, the goal is to acquire something infinite. Because space has an infinite number of points, no look-up table will suffice (see also Hay, 1974). Thus, as with language acquisition, what is acquired is a generative system, rather than specific instances. In addition, as with language acquisition, environmental input is impoverished, especially for our cases where only one or two locations are specified. Although specific issues of positive and negative evidence may not be relevant to spatial mappings, the subset rule may still be sensible. It may always be more efficient under these conditions to infer the smaller infinity so as never to have to unlearn something inferred incorrectly.

VII. Computer Mappings

The second paradigm to be described permits novel visual–motor mappings to be produced in two-dimensional space. Visual positions are presented on a standard computer monitor. Motor positions are produced using a pen attached to a tablet, which is lying flat on a table. The pen and tablet work similarly to a computer mouse: The position of the pen on the tablet is represented on the screen by a visual character. Moving the pen causes the visual character to move appropriately on the screen in real time with no noticeable delay. Figure 14 shows a rough layout of tablet and monitor. Novel mappings can be imposed via software that alters how the tablet locations are mapped onto the screen. For instance, a mirror reflection would cause the cursor to move to the left when you moved your hand to the right. Helen Cunningham (1984) was the first to document this paradigm as useful for the study of visual–motor mappings.

The first question we asked was whether two dimensions of space are independent of one another. When detecting and accommodating new visual–motor relations, do orthogonal dimensions get processed separately from one another? If they do, then it should be possible to learn about one dimension (e.g., x), without regard to what is done to the other (e.g., y). Consider two mappings, both of which create exactly the same transformation of x but with differing manipulations to y. Both mappings shrink motor space in the horizontal direction to half its former size. For instance, moving the pen left to right by 2 in. on the tablet causes the cursor to move 4 in. left to right on the screen. One of the mappings also shrinks y

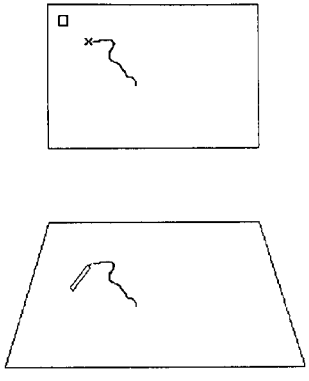

Fig. 14. Sketch of computer monitor (top) and tablet surface (bottom) used to produce computer mappings. Pen on tablet represented on monitor with an "x."

by the same amount as x, whereas the other mapping leaves y unchanged from normal.

The mappings can be characterized by very simple linear equations. The normal relation between visual position and the tablet position before transformation is: $T_x = V_x + 0$ for x, and $T_y = V_y + 0$ for y. The relation under Mappings 1 and 2 is as follows:

Mapping 1

$$T_x = 0.5V_x + 0 \qquad \text{(Minify both } x \text{ and } y\text{)}$$
$$T_y = 0.5V_y + 0$$

Mapping 2

$$T_x = 0.5V_x + 0 \qquad \text{(Minify only } x\text{)}$$
$$T_y = 1.0V_y + 0$$

If the orthogonal dimensions of space are independent of one another, then the novel x relation should be learned equally effectively in both groups because it is the same in both groups. If they are dependent, then learning about x will depend on what is done to y as well. If so, then there are two main possibilities. Shrinking just x (Mapping 2) may be easier than shrinking both x and y (Mapping 1) because on average, positions move less under that transformation. The system may be trying to minimize the amount of change. Alternatively, the system may be trying to minimize the number of geometric properties altered. This rule would make the opposite prediction, that is, Mapping 1 is easier than Mapping 2. Referring back to the hierarchy of transformations, shrinking both x and y by the same amount is less radical a transformation than shrinking only x. The former mapping (similarity) alters only the size, but the latter mapping (affine) alters shape properties as well. Mapping 1 turns a square into a small square, and Mapping 2 turns a square into a rectangle.

Thus, we can determine whether the two dimensions of space are independent, and if they are not, get some idea of how they are linked. We assigned subjects to one or the other mapping. As in the prism adaptation paradigm, normal visual–motor performance was first determined. In this paradigm, a small square target appeared on the computer monitor on each trial. Subjects moved the pen without feedback so that the cursor produced by the pen (an "X") was invisible. Subjects were told to move the pen to the position on the tablet that *would* make the X fit in the square, if it were visible, and were tested on different positions. Practice with the novel transformation was subsequently provided by making the X cursor

visible. With the distorted mapping in place, on each trial subjects tried to move the X into the small square target. There were nine different positions that together would make up a large square: three on the top row, three in the middle, and three on the bottom. Subjects practiced for a total of 9 min (135 trials), and were tested again without feedback.

The outcome was surprising and did not match any of the specific predictions. When the mapping minified both x and y, both x and y did shrink. That is, subjects did use a smaller range of the tablet for both x and y directions than they did before practice. When the mapping minified only x, x did shrink—and by an amount comparable to the other mapping. However, the y direction also became smaller, even though it should not have! The data are shown in Fig. 15 in which it can be seen that basically a square turned into a smaller square for both groups, even though for one it should have been turned into a particular rectangle. Although the results were unexpected, they do support the view that preserving geometric properties may be important. Shape was preserved, even when it meant imposing a transformation that was not correct. It is apparently easier to rescale both x and y equally, then just to minify x. The slope parameters describing the data are as follows:

Mapping 1 (Minify both x and y)

Pretransformation	Posttransformation
$T_x = 1.13V_x$	$T_x = 0.80V_x$
$T_y = 1.07V_y$	$T_y = 0.75V_y$

Mapping 2 (Minify only x)

Pretransformation	Posttransformation
$T_x = 1.17V_x$	$T_x = 0.83V_x$
$T_y = 1.12V_y$	$T_y = 0.85V_y$

But did y really shrink because of x, or would it have happened anyway? Note that before any transformation, a larger area of the tablet was used than should have been (pretransformation slopes start out greater than 1.0). Perhaps the range would shrink anyway following any kind of practice with the equipment. Figure 16 (top) shows that this is not the case. When an equal amount of practice is given with the normal mapping instead of a novel mapping, the y range of the tablet (as well as the x range) does not shrink. In fact, initially the range expands even more. This may occur because the tablet is larger than the computer screen and there may

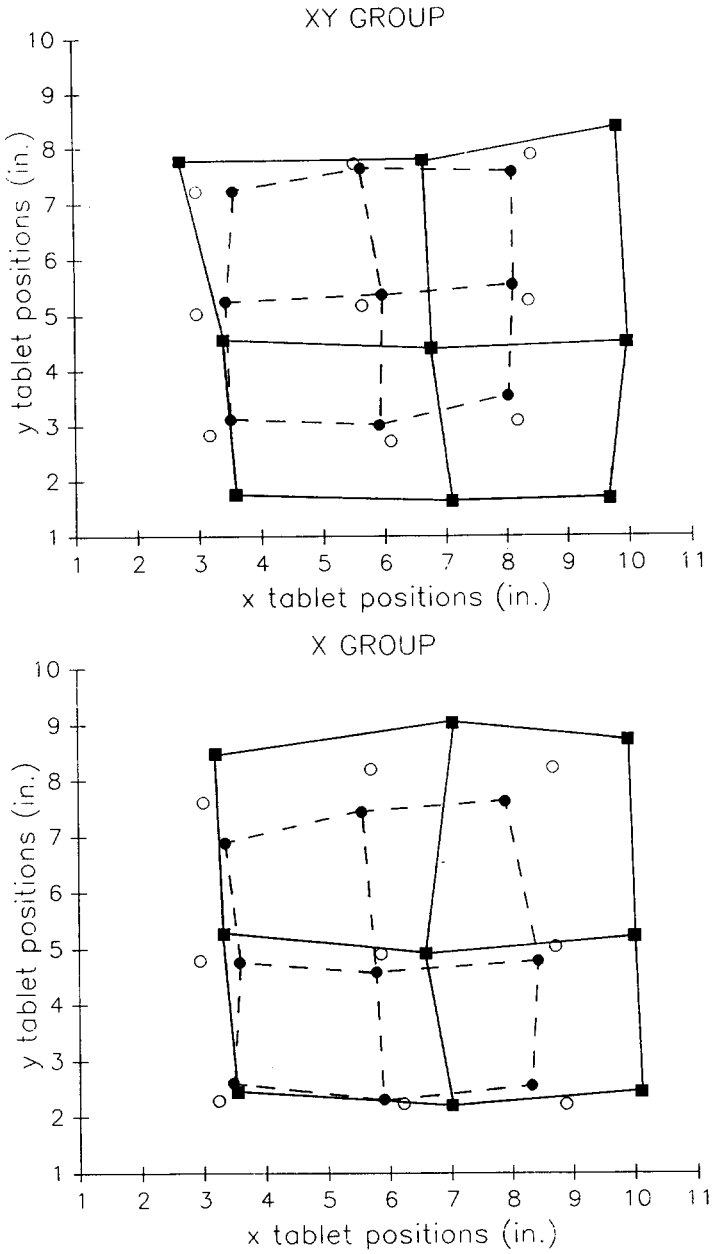

Fig. 15. Tablet positions before and after training for the "independence experiment." Connected squares show positions before training, unconnected (unfilled) circles after 6 min. of training, and connected (filled) circles after 9 min. Lower left corner of tablet is located at (0,0). "*XY* group" is the mapping that shrinks both *x* and *y*; "*X* group" is the mapping that shrinks only *x*.

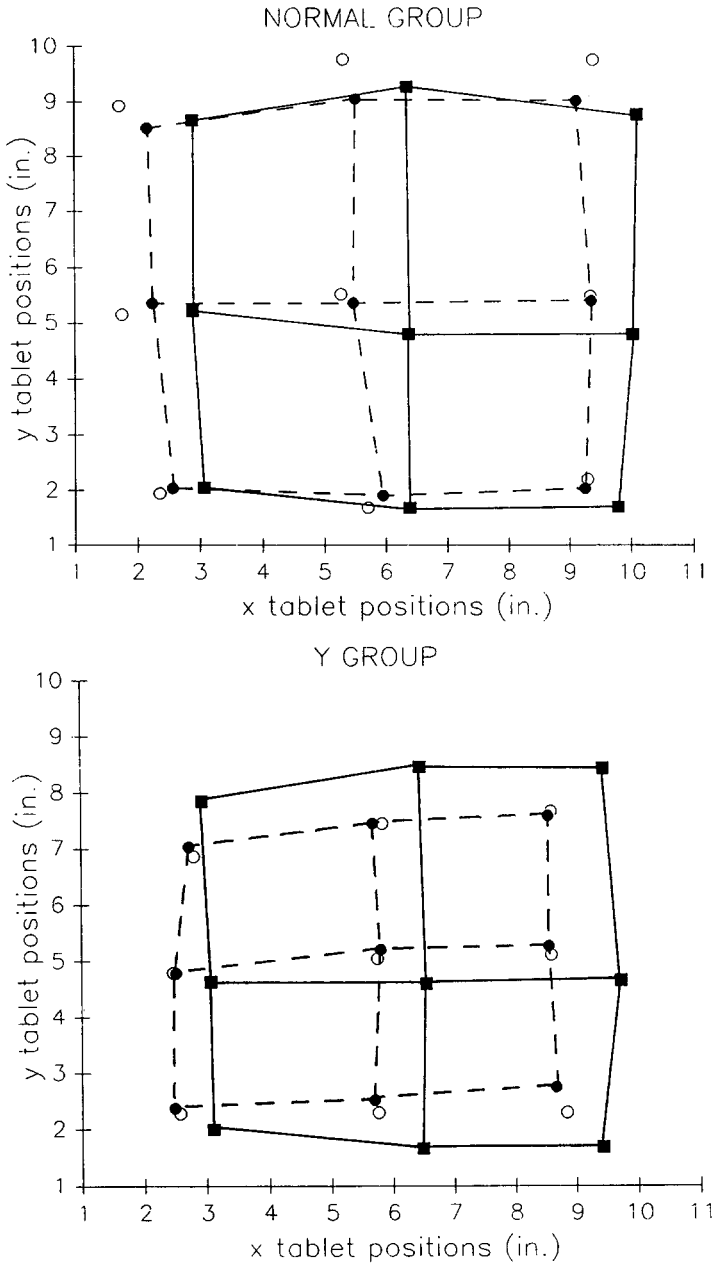

Fig. 16. Tablet positions before and after training for the "independence experiment." Connected squares show positions before training, unconnected (unfilled) circles after 6 min. of training, and connected (filled) circles after 9 min. Lower left corner of tablet is located at (0,0). "Normal Group" is the mapping in which neither x nor y is changed; "Y group" is the mapping that shrinks y.

be a natural tendency to map the whole tablet to the whole screen. A tendency to use a larger tablet range than the normal mapping dictates can account for the pretransformation values as well.

But there is at least one puzzle. To complete the fourth permutation, we ran another group where y was shrunk, but not x. Figure 16 (bottom) shows that y does not pull x with it as obviously as x influences y. The asymmetrical influence may be due to the asymmetry imposed by the paradigm. The x dimension maps straightforwardly from the tablet to the screen, but the y dimension does not. Moving the hand from left to right on the tablet causes the cursor to move from left to right on the screen. Moving the hand in a forward direction (close to far) on the tablet, which lies flat on the table, causes the cursor to move from the bottom to the top of the screen. There is evidence from research on interpreting "you are here" and other maps of the environment, that people naturally assume *forward* corresponds to *up* (Levine, 1982). Yet perhaps subjects are not completely certain of that correspondence and may not know how to align those two distinct spatial dimensions. In this paradigm, the y dimension may consequently be more labile than x in the sense that it takes guidance from the more certain dimension, especially given novel transformations. One of the many intriguing questions that remain to be investigated is whether x and y will be dependent or independent, and whether they will be symmetrical or asymmetrical, when the y mapping from tablet to screen is made more straightforward.

Having determined that at least under some conditions x and y are dependent, and that shape-preserving transformations seem preferred, we still wanted to test whether the hierarchy of transformation geometry is a useful way to categorize mappings. The following experiment should be viewed as exploratory, with both more data analyses and experiments needed. Five different mappings corresponding to the five different levels of the hierarchy were given to the same subjects on different days and in different orders. The particular mappings chosen are as follows. The isometric mapping shifted motor space down with respect to visual space by 1.5 in. That is, the position of the pen on the tablet is lower than where it appears on the screen, for all positions. The similarity mapping uniformly shrunk both x and y to half their normal range, as in the last experiment. The affine transformation sheared space by essentially "pulling" x axes diagonally so that the y position on the tablet depends on the x position on the screen. For instance, to get a path on the screen to be perfectly horizontal from left to right requires moving the pen diagonally from lower left to upper right. A square region of the screen is turned into a sheared square on the tablet. The projective transformation turned a square region into a trapezoid. And finally, the topological transformation we chose

mapped straight lines of space to curved lines and turned a square into a small circle. The transformations all moved the upper left target down by exactly the same amount (1.5 in.). The left column of Fig. 17 shows how the transformations should alter behavior to the nine test targets if learning were perfect; the right side of Fig. 17 shows what actually was learned for each transformation.

Overall, the results are promising. The more geometric properties altered by the transformation, the harder the mapping was to acquire. The first three levels can be compared relatively directly by fitting linear equations, because all three are linear mappings in two dimensions. Figure 18 (top) shows the change in the parameter(s) that should have changed, as a percentage of optimal amount of change for that transformation. The Y displacement mapping was learned 100% (change in y intercept); next was the minification mapping at an average for both parameters of 70% (change in x slope, change in y slope), followed by shear at 26% (change in y as a function of x).

To compare all the groups, the positions for perfect acquisition for each mapping were determined, followed by a calculation of how far the pen position was from the optimal position for each target (average square root of x^2 and y^2). The average distance from the optimal positions following training are shown in Fig. 18 (bottom). The Y displacement group is closest to the destination, followed by circle, minification, shear, and trapezoid. The circle group appears better accommodated than was predicted from the hierarchy; however, this result comes from accommodating the minification part of the transformation, rather than from altering the shape. There is little evidence that the straight lines of the square became curved lines of a circle.

Although promising, the data are not conclusive. The distance measure is problematic because it does not take into account the shape of the mappings, and could be misleading. There are different possible interpretations of the outcome of the experiment. For instance, the data could support a view that transformations of scale are easy, and everything else difficult. The geometric hierarchy as well as other possibilities are being explored.

Is this perceptual learning? It is difficult to determine whether the constraint that an object cannot be in two places at once, or any other constraint, has been violated in the pen and tablet paradigm. Without information that there is an internal error, the motivation for perceptual learning would be absent. If we consider again the prism adaptation paradigm, the error is suggested by feeling an object (the hand) in one place but seeing the object (the hand) in a different place. Welch (1972; 1978; 1986, p. 10; Welch & Warren, 1980) points out that the object localized by the two modalities

must be judged as the *exact same object*. Otherwise, there is no discrepancy between the modalities. If you feel your hand straight ahead of your body, but assume that the visual image you see a few inches away is someone else's hand, then there is no motivation for adaptation. The information wouldn't be in violation of any constraint, because different objects can be in different places (see also Bedford, 1992a).

Returning to the digitizing tablet, are there sufficient cues present to assume that the information obtained visually and the information obtained through the hand refer to the same object? The motor information is the feel of the hand's (and pen's) position and motion on the tablet, yet the visual information isn't of a hand, or part of a hand, or the pen. It's the character "X", an abstract representation. This is also to a large extent true of our prism adaptation paradigm, where only a small, round LED on the hand was visible. However, in the digitizing tablet paradigm, the hand/pen and its visual representation are also located very far from one another, which could decrease the same object assumption. The paradigm imposes three spatial transformations of its own, even if the normal mapping is used: The tablet is lower down than the screen (a y displacement), it's closer to the observer (a z displacement), and it's flat on a table, whereas the screen is upright (rotation around an x axis). Large lateral displacements in prism adaptation produce less percentage adaptation, which could be due in part to some object assumption breaking down. If you look at your hand through a prism that displaces the image 10°, your hand temporarily feels as if it's located where you see it (visual capture). If you look at your hand through a 30° prism, visual capture breaks down and your hand looks and feels as if it's in two different positions.

Yet there could be other factors that compensate for the large differences in spatial position produced by the digitizing table paradigm. Held and Durlach (1989) suggest that motion plays a major role in object identity. The visual- and motor-motion information are highly correlated, with the movements occurring at the same time and with the same velocities and accelerations. The bottom line is that exact circumstances under which two distinct samples will be judged, by the perceptual system, to refer to the same object are not precisely known.

Another approach would be to see whether the change in aiming for the target is due to a genuine perceptual change, in which either the hand feels as if it's in a different position on the tablet, or the target on the screen looks as if it's in a different position. We asked subjects after each experiment to describe any transformation that they were aware of, although not in those words. Some of the subjects report intentionally using a strategy to get the "X" in the square—such as moving their hands a smaller amount on the tablet in the minification of the x and y conditions—and

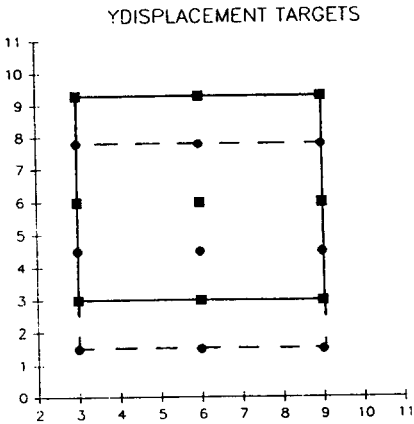
YDISPLACEMENT TARGETS

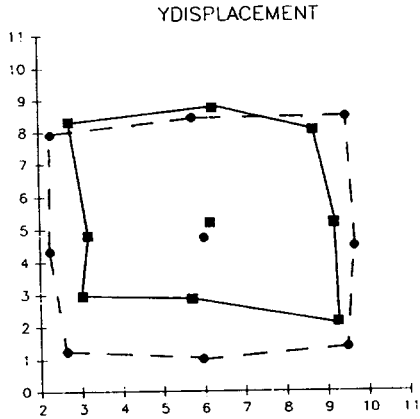
YDISPLACEMENT

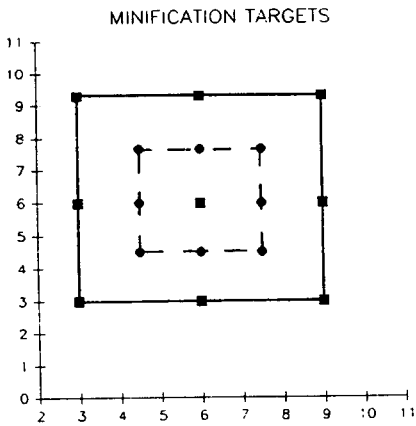
MINIFICATION TARGETS

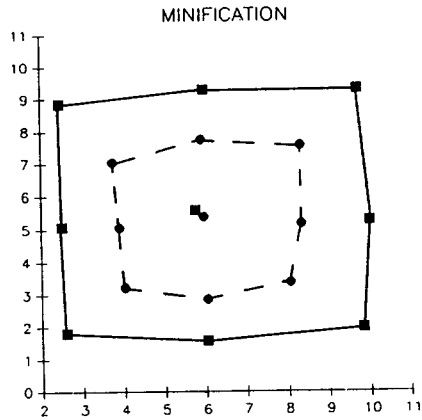
MINIFICATION

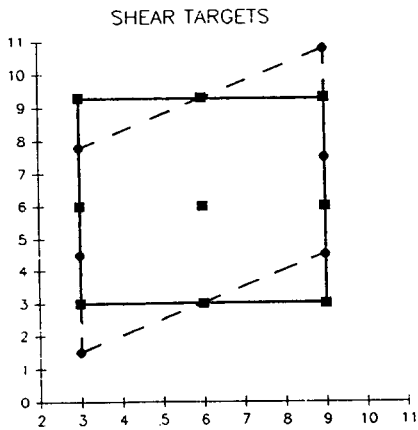
SHEAR TARGETS

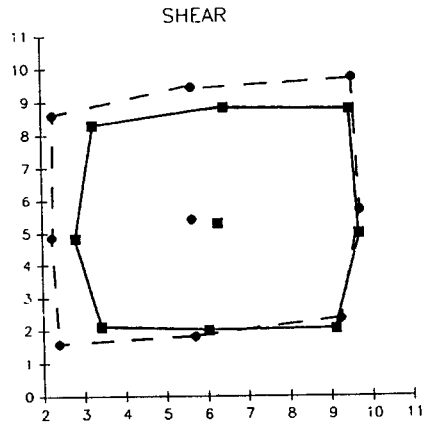
SHEAR

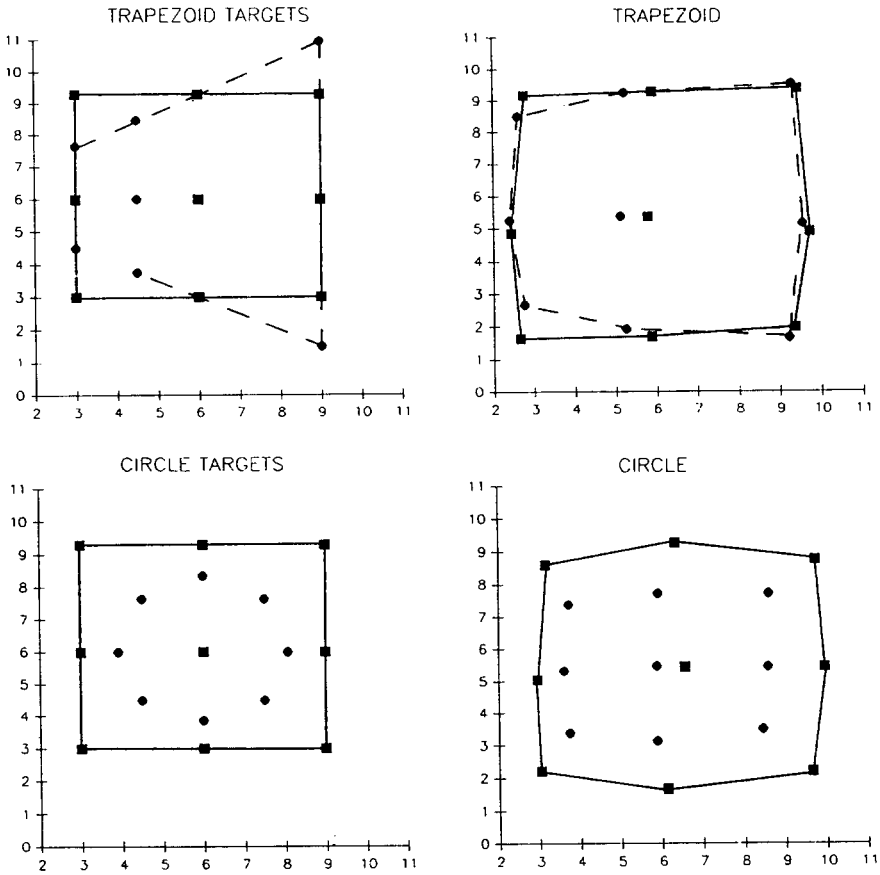

Fig. 17. *Continued.*

using that strategy for testing as well. Interestingly, so far the data for those subjects appear indistinguishable from the others. Yet the presence or absence of conscious awareness does not necessarily map neatly onto cognitive and perceptual learning anyway. Direct tests of perceptual changes, such as transfer to entirely visual tasks or entirely motor tasks, are being investigated.

Note that in any event, the paradigm remains of interest. The ease with which people use computer mice and learn video games may suggest that

Fig. 17. Geometry experiment. Left column shows the transformation of space required of subjects for the nine test targets. Right column shows the outcome. Data shown before training (connected squares) and after 6 min of training (connected circles). The circle data are not connected.

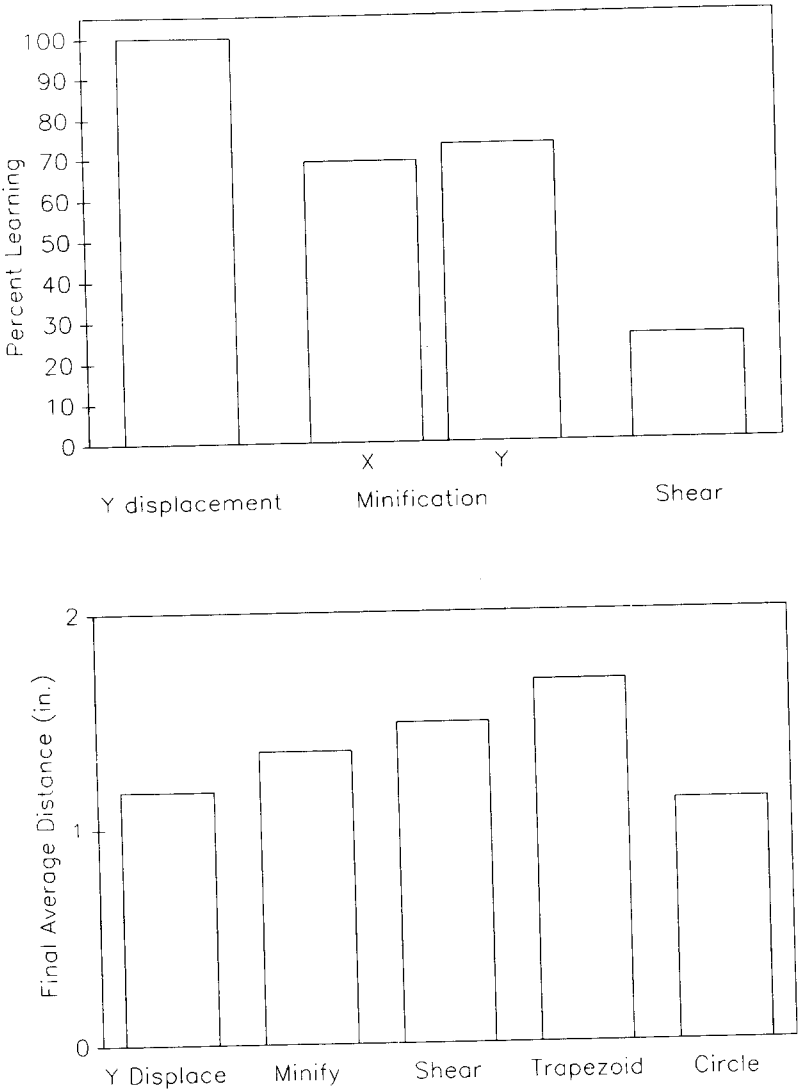

Fig. 18. Geometry experiment. Top panel shows the percentage change in the relevant parameter(s) for the three linear transformations. Bottom panel shows the average distance from the optimal positions following training for all five transformations. (See p. 40.)

the rules of perceptual–motor coordination and its maintenance are accessible for purposes other than to correct sensory systems. Tasks with similar characteristics—such as needing mappings between entire dimensions, and referring to space—may be "parasitic" on perceptual learning

proper, even when an internal error has not, strictly speaking, been detected. This view may be similar to the view of "accessibility" of adaptive specializations (Rozin, 1976), in which mechanisms evolved for one purpose can be exploited for other purposes. Although these ideas may weaken the modular story of perceptual learning to some extent, they may nonetheless prove true.

Adaptation to the rearrangements of space and to learning new mappings between spatial dimensions appear to differ in significant ways from apprehending new information about the world. Is adaptation unique, or do other processes differ from world learning also? I turn now to a different phenomenon, which has recently been gaining in interest.

VIII. The McCollough Effect

An example of plasticity entirely within the visual system is a phenomenon known as the McCollough effect or the orientation-contingent color aftereffect, discovered in 1965 by Celeste McCollough. A vertical grid consisting of alternating black and magenta bars is shown for a few seconds, followed by a grid of horizontal bars colored black and green (see Fig. 19). The two patterns are alternated repeatedly for several minutes. This experience leads to the perception of colors that are contingent on the orientation of a stimulus. The white portion of black and white vertical bars look slightly green, and white horizontal bars look pink.

The explanation for this phenomenon has proved elusive, candidates ranging all the way from fatigue within simple neural detectors to a Pavlovian conditioning interpretation (e.g., Dodwell & Humphrey, 1990; Harris, 1980; McCollough, 1965; Murch, 1976; Seigel & Allan, 1992). I suggest first that the phenomenon is properly thought of as an instance of *learning* as described in the first section. It is a change from experience that evolved to make us better. Precisely what function it serves will be elaborated

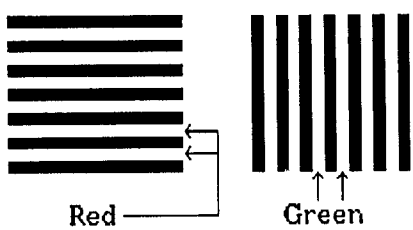

Fig. 19. Stimuli for a McCollough effect experiment.

later. Second, within the category of learning, it belongs within the division of internal perceptual learning, rather than that of apprehending new information about the world. This point can be illustrated by an experiment conducted recently in my laboratory, showing that the mechanism responsible for the contingent aftereffect cannot make use of properties in the world. In the standard McCollough effect, tilting one's head to the side by 90° causes the illusory colors seen on the test patterns to reverse (Ellis, 1976; McCollough, 1965). For example, if vertical lines appear pink and horizontal lines green, then tilting the head makes vertical lines appear green and the horizontal lines pink. Yet when you tilt your head, objectively vertical lines continue to be perceived as vertical even though the retinal orientation becomes horizontal. That is, the illusory colors appear yoked to the retinal orientation, rather than to the perceived or objective orientation of a stimulus.

Karen Reinke and I (Bedford & Reinke, 1992) pursued this finding by asking whether color *could* be made contingent on the objective or perceived orientation. When inducing the standard McCollough effect, color is paired both with retinal and with objective/perceived orientation. Logically, acquiring either dependence will correctly reflect the contingencies of the environment. It is possible that when both contingencies are made available, the retinal properties overshadow the regularity involving the real orientation. Under these artificial circumstances, learning about the world may be prevented. In the real world, it can be argued, the continual motion of an observer may prevent any longterm retinal correlation. That is, it could be that the purpose of the mechanism reflected by the McCollough effect is to apprehend new features of the environment—that vertical lines and red color go together in some sense. We asked, if subjects received only pairings between color and objective orientation, would they learn the connection?

To answer this question, we dissociated retinal from perceived coordinates during training by having each subject view the stimuli with her head tilted 90° for half of the trials. Consider a red grating, which is objectively vertical, alternating with a green grating, objectively horizontal. If viewed when upright, red is paired with both actual vertical and with retinal vertical, and green with both actual and retinal horizontal. If those same stimuli are next viewed with the head tilted, red remains paired with actual vertical and green with actual horizontal, but the retinal relation has reversed. Red is now paired with retinal horizontal, and green with retinal vertical. Equal numbers of trials when the head is upright and when it is tilted produce a zero retinal contingency where both retinal horizontal and vertical lines will each be red half the time and green the other half. The colors remain perfectly correlated with the actual orientation of the stimuli

(red-vertical/green-horizontal). We found that very little was learned in this condition compared both to a standard condition in which both contingencies are present, and to a condition that left only the retinal contingency. It appears to be extremely difficult to get the illusory color under the control of the actual orientation of a stimulus.

We use this finding as a plausibility argument to suggest that some classes of explanation for the McCollough effect will not be very useful. Any account for which the goal is to apprehend new information about the world would be ill served by a process that is insensitive to that information. We suggest instead a theory that is based heavily on perceptual learning as distinct from other forms of learning.

The theory has two assumptions, the first a condition necessary to produce learning and the second a claim about the content of what is learned (Bedford, 1992a). The assumptions are first summarized and then followed by a more detailed explanation.

1. Objects should not change their color when the head (or object) is tilted. If they appear to do so, then an internal correction is necessary the manifestation of which we call the McCollough effect.
2. The McCollough effect involves a mapping between entire dimensions of orientation and color rather than learning two specific red-vertical and green-horizontal associations.

According to the first assumption, when an experimental induction procedure pairs a vertical grating with red, and a horizontal grating with green, it is actually providing the visual system with the information that as the retinal orientation of an object—the grating—changes, so does its color. This is problematic because when the orientation of an object on the retina changes, whether due to head tilt or object tilt, color should not. Objects do not change their color when the head (or object) is tilted. If they appear to do so, an internal correction is necessary, the manifestation of which we call the McCollough effect. There will be an internal correction for the perceived color differences, so that different orientations of the same object will no longer have different colors. Red vertical lines and green horizontal lines should eventually appear as exactly the same color if learning is perfect. If tested on white lines, vertical will appear green and horizontal will look red.

Unless there is some information to inform the visual system that it is malfunctioning, perceptual learning will not occur. I suggest that we have a constraint that an object does not change its "color" over time, analogous to the constraint that an object cannot be in two places at the same time. If information contained in the proximal stimulus does not contradict the constraint, then internal changes are not necessary and will not occur.

Consequently the co-occurrence of colors and orientations is problematic, from the viewpoint of the visual system, only if the different colors refer to the *same* object in different orientations. Different objects can have different colors. This assumption makes strong predictions about when contingent aftereffects will and will not occur, an issue that has been the subject of continuing debate (e.g., Dodwell & Humphrey, 1990; Harris, 1980; Seigel, Allan & Eissenberg, 1992; Skowbo, 1984).

The extent to which an induction procedure is successful will depend on the extent to which the perceptual system is convinced that the same object is involved. In the standard phenomenon, red and green are paired with vertical and horizontal gratings, which leads to illusory colors contingent on orientation. Also successful have been red and green paired with a single vertical and single horizontal bar, a triangle pointing up and a triangle pointing down, (vertical) wide stripes and (vertical) narrow stripes, two concentric circle patterns of different spatial frequencies, two different directions of motion, two different velocities, two disks of different lightnesses, and two different lightnesses of a surrounding frame (Breitmeyer & Cooper, 1972; Harris, 1980; Hepler, 1968; Lovegrove & Over, 1972; Mayhew & Anstis, 1972; Siegel, Allan & Eissenberg, 1992; see Stromeyer, 1978, for a review). All these stimuli lead to color contingent on the appropriate feature.

In the present interpretation, these successes are sensible because in all cases the two members of the pair can be interpreted as referring to a single object. Two concentric-circle patterns of different spatial frequencies can refer to the exact same concentric-circle pattern seen at different distances, because spatial frequency changes as a function of distance. A vertical bar and a horizontal bar can refer to the exact same bar viewed either with the head upright, or with the head tilted by 90°. A light-gray disk and a dark-gray disk can be the same disk under different overall levels of illumination. Gratings moving up and moving down on the retina can result from the exact same grating viewed with different up and down head motions. Different velocities can result from different speeds of head motion. Once the "same-object" constraint is met, then one of the required conditions for a discrepancy is met. When the retinal image of a single object tilts, its color should not; when the distance between an object and an observer changes, the object's color should not; when the distance from an object changes, the object's velocity should not, and so on.

Stimuli that do not succeed are those not readily interpretable as different states of the same object. Concentric circles of one color alternated with radiating lines of another did not produce a color effect. Pairing one direction of motion with small disks, and the other with small triangles also

did not produce a contingent aftereffect, nor did alternating a green cross and a red square (Fidell, 1968, reported in Skowbo, Timney, Gentry & Morant, 1975; Foreit & Ambler, 1978; Mayhew & Anstis, 1972). There is a bias toward reporting successes, but a number of researchers have concluded that arbitrarily chosen stimuli tend not to work (e.g., Dodwell & Humphrey, 1990; Skowbo, 1984). As White and Riggs (1974) note, no one has ever reported color contingent on different makes of automobiles. From the viewpoint of the visual system, none of these results are problematic. For instance, the square and cross were ineffective because the retinal image of a cross and the retinal image of a square cannot refer to the same object—assuming objects maintain rigidity. Pairing a red square and a green cross requires no action, because different objects can be different colors. The same argument holds for pairing small triangles and small circles. There is no normal transformation that turns triangles into circles. That is, retinal images of triangles and circles always refer to different objects. No internal malfunction will be detected, or corrected, and no illusory colors manifested.

The above is a simplified version for sake of clarity, dividing conditions into simply whether they will work or won't work. It is likely that the distinction is one of more or less effectiveness, rather than of all or none. For instance, it is probably more accurate to say that the retinal image of a circle and of a triangle is *less likely* to refer to the same object than a small circle and a large circle. Non-natural transformations are apparently sometimes acceptable for object identity (see Shepard, 1984, pg. 430). This probably explains why a (red) square and a (green) cross are ineffective when using the same procedure that is sufficient for oriented lines and color (Foreit & Ambler, 1978), but why the aftereffect can be detected when using a test procedure designed for measuring very small effects (Siegel, Allan & Eissenberg, 1992). As noted earlier in the section on computer mappings, the circumstances under which two distinct instances are judged by the perceptual system to refer to the same object are not precisely known. It should be possible to use the criteria of object identity uncovered in other areas, such as apparent motion (e.g., Chen, 1985; Warren, W. H., 1977), certain types of priming (see Kahneman, Treisman & Gibbs, 1992), and unity in infants (Spelke, 1990) to predict how effective stimuli will be in the McCollough effect.

Stimuli can fail to suggest a malfunction in ways other than by referring to different objects. The two stimuli can refer to the same object, but not undergo a change in color. For example, the alternation of a vertical grid colored green and a horizontal grid also colored the same color green does not induce any contingent aftereffect (Humphrey, Dodwell & Emerson, 1985). In the current interpretation this result is sensible because different

retinal orientation of the same object did not change color. There is no discrepancy here, just a green object.

Other learning-based accounts have had difficulty explaining and predicting which stimuli can and cannot induce contingent aftereffects. The most well developed traditional learning account maintains that the McCollough effect is an instance of Pavlovian conditioning (e.g., Allan & Siegel, 1986; Murch, 1976; Siegel & Allan, 1992; Westbrook & Harrison, 1984). A naive instantiation of this view predicts that any arbitrarily chosen stimuli can be associated, an account clearly falsified by the data. A more sophisticated view argues that based on phenomena of selective association in animal learning (e.g., taste aversions), not all stimuli have to be equally associable (see, e.g., Harris, 1980). This would allow for failures as well as successes. However, it does not allow one to specify *which* stimuli will work, unlike the present interpretation (see Bedford, 1992 for additional problems).

Once the circumstances that induce perceptual learning are present, something gets learned. The question becomes, What is the content of that learning? The second assumption is that part of what is acquired is a mapping between the dimensions of orientation and of opponent red/green color (cf., Hurvich & Jameson, 1957). This suggestion is in contrast to associative models, which assume that what is learned is the specific red-vertical and/or green-horizontal pairs presented during induction. Models based on independent connections between individual stimuli make different predictions from those based on connecting entire dimensions, including whether one can distinguish the trained stimuli from the untrained stimuli in performance and whether individually trained pairs would combine additively to produce all possible combinations of trained pairs (see Bedford, in press; Koh & Meyer, 1990). I focus here on one particular difference, the information necessary for a correlation, because there is relevant data from the McCollough effect and because the interpretation of the data has been controversial.

The successful correlation of two individual stimuli, such as red and vertical, requires integrating information about the presence and absence of one of those stimuli with the presence and absence of the other. Figure 20 shows the four possible combinations along with different outcomes. For instance, if S_1 and S_2 co-occur often, and S_1 does not occur without S_2, nor S_2 without S_1, then the correlation between S_1 and S_2 is high—S_1 and S_2 appear and disappear together. If S_1 and S_2 co-occur often, but S_1 also occurs often without S_2, and S_2 without S_1, then the correlation is low—S_1 and S_2 are independent and the co-occurrences were coincidental. Thus, a mechanism correlating two individual stimuli would have as input the information in those cells.

S_1 \ S_2	Present	Absent
Present	100	?
Absent	?	100

S_1 \ S_2	Present	Absent
Present	100	0
Absent	0	100

High (Perfect) Correlation

S_1 \ S_2	Present	Absent
Present	100	100
Absent	100	100

Low (0) Correlation

Fig. 20. High and low correlation between two stimuli (S_1 and S_2) with number of pairings between S_1 and S_2 held constant (100).

If instead a mechanism correlates the values along two entire dimensions, it is not obtaining information showing that the dimensions appear and disappear together, but rather that *if* the two stimulus dimensions co-occur, then the values of one dimension may (high correlation) or may not (low correlation) be systematically related to the values along another (in a linear fashion for simple linear regression). The relevant information does not come from the four cells described above, but from all pairs of values from the two dimensions. For instance, the correlation of height and weight takes as its input pairs of height-weight values. If the weight of a 6-ft person was unknown, it would not affect the correlation, whereas if the weight of a 6-ft person was 90 it would[5].

Now we can consider the controversial finding from the McCollough effect. If trials of a red homogeneous patch are interspersed among red-vertical and green-horizontal induction trials, they do not reduce the magnitude of the contingent aftereffect that results (Skowbo & Forster, 1983). This is problematic for a single stimulus account because, as described above, this manipulation decreases the correlation between the two stimuli. However, the finding is well accommodated within a dimension account. The presence of red without vertical is simply a missing data point,

[5] The correlation of two individual stimuli can be viewed as a simplified subset of correlation proper, where the "dimension" for an individual stimulus consists only of two values—presence and absence.

analogous to entering height without weight into a correlator. It also accounts for why interspersing, instead, achromatic vertical lines (i.e., vertical lines without red) does reduce learning. This is not a missing data point because achromatic vertical lines do pair orientation with color; *vertical, white* will be entered into the correlation.

The dimension approach suggests manipulations that largely have not been conducted. For instance, mappings between dimensions shift the focus from simply whether you get a connection, to the type of connection. Must orientation and red/green color always be related linearly? If we pair vertical lines with green, and horizontal lines with a less saturated green rather than red, will the illusory colors seen be related linearly to orientation? How about the presentation of multiple orientation–red/green color pairs rather than only two pairs? Will linear functions always be fit even if the relation presented is nonlinear? Are some mappings more preferred and more easily learned than others? It is easy to see why research has been limited to only two pairs. Because only two points are needed to determine a line, two pairs are sufficient to recalibrate the entire dimension. Under these circumstances, it can appear as if one or two specific associations are the core phenomenon, rather than reflecting a degraded version of the actual phenomenon of mappings between entire dimensions. Note that with rich internal constraints on this learning system, even a single pair can be sufficient to drive a change throughout the dimension, as in prism adaptation (see Bedford, 1992a).

In addition to dimensions, another answer to the question of the content of learning should involve which system changes, an often ignored aspect of the McCollough effect. Why does color become contingent on orientation and not vice versa? In fact, that relation can also be learned (Held & Shattuck, 1971) but the factors that lead to one or the other have not been systematically explored. The perceptual learning view approaches this question by asking which system is malfunctioning. If an error is detected, correct resolution requires first inferring the source of the error. Does the color dimension need recalibration because it is erroneously varying with changes in retinal orientation (head tilt)? Or is it the orientation dimension that needs recalibrating, because it is erroneously varying with changes in relative luminance? A related problem that must be solved is whether color is erroneously varying with orientation (if that direction is inferred), or with some other dimension. Research that can be viewed as related involves effects labeled *blocking,* in which past experience can bias decisions about the source of the error (e.g., Brand, Holding, & Jones, 1987; Sloane, Ost, Etheridge, & Henderlite, 1989; Westbrook, & Harrison, 1984; Siegel & Allan, 1985). Framing the problem this way should also enable selective control over which system changes by trying to bias one or the other interpretation.

Having elaborated briefly on the two assumptions, I will discuss a few related general questions to keep in mind. Why does it make evolutionary sense to have these internal corrections? We need to know how such a mechanism could have arrived there in the first place, and why it would be adaptive in an evolutionary sense of the word to have such a mechanism. Clearly at an abstract level, it is advantageous to have optimally functioning sensory systems, without which fatal errors are more likely to result. But we need more specific hypotheses for the specific internal corrections. In the case of the McCollough effect, why have a mechanism prepared to correct for observed color changes of an object seen in different orientations, if an object does not change its selective reflectance of wavelength?

I suggest that the function is an extension of maintaining perceptual constancy. *Constancy* refers to perception of the property of an object remaining constant despite continually varying input. For instance, size constancy refers to the perception of the size of an object remaining constant, despite changes in retinal size with distance. Color constancy refers to the perception of the color of an object remaining constant despite changes in relative luminance that result from different lighting conditions. The point in both instances is that the size and "color" of an object in the world typically do not change. Consequently, it is advantageous to perceive them as non-changing, so as to represent the world accurately. It would defeat the purpose of constancy, if sometimes the color on an object erroneously appeared to change with head tilt. It may therefore be useful to ensure that when you tilt your head, not only shouldn't perceived orientation change ("primary constancy"), but neither should perceived size or color or motion, and so forth ("secondary constancy"). But why prepare for that possibility if it can never happen? Is nature paranoid enough to prepare for experimenter intervention?

The missing piece of the puzzle may be provided by Held (1980), who points out that there are natural situations in which the color can indeed appear to change when viewed from different orientations. There is a common optic defect in which the optical axis of the lens is not aligned with the fixation axis. As the "ocular prism" moves with the eye, color fringes at edges will appear that are contingent on the orientation of the stimulus with respect to the observer. Thus, there appears to be a naturally occurring situation that would create the otherwise impossible covariation between orientation and color. Consequently, it is sensible that a corrective mechanism already exists to handle the natural discrepancy, which can also be tapped by experimenters who artificially induce a similar discrepancy. We refer to these mechanisms as maintaining *secondary constancy*, to distinguish them from classic perceptual constancy, which we can call *primary constancy*. Due to the importance of keeping perception constant

and veridical while besieged with bodily movements and changes in lighting conditions, the evolution of corrective mechanisms would be adaptive to the organism.

What happens if the motivation for perceptual learning is absent? The present view does not claim that nothing will be learned in some situations where stimuli co-occur. For instance, if a triangle is repeatedly red and a circle repeatedly green, an organism would be foolish not to take notice of the pervasive contingency. However, what will be learned will be about the external world and not about internal perceptual systems. Such learning could manifest itself as changes in thoughts, expectations, beliefs, or behaviors, but there is no need to alter what is seen because there is no internal malfunction.

IX. Conclusion

The general direction of this chapter has been to argue that *perceptual learning* is a distinct type of learning, with its own inputs, internal states, outputs, and function. In addition, the issue of how this category is related to other changes that result from experience was briefly discussed in the beginning.

Considering the issue of internal states, or content of learning, one important idea involves dimensions. Perceptual learning processes may be constrained to operate on entire dimensions, rather than on individual stimuli along those dimensions. Adaptation to rearrangements of space involves rules about entire dimensions rather than rules simply about individual associations between discrete locations of space. In addition, I suggested that study of the McCollough effect might benefit from viewing what is learned as a connection between entire dimensions of orientation and opponent color, rather than as the formation of specific associations such as *vertical* and *magenta*. The emphasis on dimensions for perceptual learning may distinguish itself from world learning, which does not have the same limitations.

It is sensible that world learning not be restricted to dimensions. Objects that are perceptually similar do not necessarily have the same consequences; consider, for instance, cats and mountain lions. The world acts on individual objects and events, and the learning mechanisms should reflect this. Consequently, mechanisms that apprehend new information about the world should be equipped to handle individual stimuli as independent. On the other hand, perceptual learning may be subject to different pressures. These processes keep sensory systems functioning optimally, and sensory systems may instead be built upon entire stimulus

continua. Some fundamental dimensions both at the psychological and physiological levels appear to be spatial frequency, spatial location, orientation, and distance (e.g., Graham, 1992). Entire systems get recalibrated when learning is needed.

Other perceptual learning phenomena may involve dimensions as well. One example is the "criterion shift rule" (Warren, R. M., 1985) or "adaptation level" (Helson, 1964), not to be confused with adaptation to rearrangements of space. This set of effects involves shifting the interpretation of an entire stimulus dimension as a result of experience with a few instances along the dimension. For instance, a temperature of 65°F feels cool in Tucson, Arizona, because of the usually higher temperatures, but feels warm in Montana because of the usually lower temperatures. A given physical weight will feel heavier or lighter as a function of the weight of objects experienced previously. These effects differ from adaptation to spatial rearrangements discussed in this chapter. The criterion shift rule does not refer to detection of new relations between pairs of dimensions, but instead to the reception of information from a single dimension. Yet, like all perceptual learning phenomena, its purpose is to keep perceptual systems functioning optimally (Warren, R. M., 1985). In addition, it appears to operate on dimensions.

Another phenomenon that may involve dimensions is the maintenance of an accurate internal clock, known as the *entrainment of circadian rhythms*. The process is not usually regarded as perceptual learning. However, entrainment may be better thought of as perceptual learning than as the representation of information from the world. Our internal sense of the time of day is kept correct by input from daylight from the environment. The function of the environmental information is not to create the underlying representation of time—that already exists. Instead, it is used to check whether the clock is accurate, similar to other perceptual learning systems that make use of environmental input to check accuracy. For the present discussion, one critical aspect is that entrainment appears to operate on dimensions also. Limited input from the environment at one time of day is sufficient to reset the clock for the remaining hours of a 24-hr day.

The view of dimensions as important to perceptual learning suggests that whenever a phenomenon appears to be perceptual learning, we should investigate whether the underlying mechanism acts on dimensions, and in addition, that whenever a mechanism appears to operate on dimensions, we should investigate whether it is really an instance of perceptual learning. Whether all perceptual learning phenomena involve dimensions, or whether that property identifies one particular subclass of perceptual learning, is a potentially important question that remains to be answered.

Considering the issue of input, or circumstances that produce learning, perceptual learning first requires the inference that internal sensory systems are functioning suboptimally. This condition is not required for world learning. One way the inference can occur is if new information obtained via vision, audition, and so forth is discrepant with internal assumptions about real-world objects. The assumption challenged in a prism-adaptation experiment is that an object occupies only one place at one time, and for a McCollough effect experiment that an object stays the same color over time. Both of these laboratory phenomena tap into mechanisms needed in the real world for accurate spatial localization and color constancy, respectively.

In addition, verifying that a constraint has been seemingly violated may involve a decision about object identity. In prism adaptation, what you are seeing and what you are touching may or may not refer to the same object. In the McCollough effect, the two samples obtained at different times (e.g., vertical grid and horizontal grid) may or may not refer to the same object. For both phenomena, if the two samples are judged to refer to the same object, then a problem will be detected, and if not, then there is no evidence of internal malfunctions. Different objects can be in different places and can have different colors. The decision about when two distinct samples refer to the same object, whether they are from two different modalities at the same time, or at two different times from the same modality, is critical both for perception and perceptual learning. The study of perceptual learning will benefit from a good understanding of how this decision is made.

If we consider the issues of output, and of function, perceptual learning is manifested by actual changes in how the world is perceived under identical physical stimulation. Perceptions are modified, which preserves internal constraints about objects and restores the systems to what is believed to be good working order. Perceptual learning processes keep perception accurate to allow world learning to occur.

The field of cognition has started to recognize that learning processes must play a central role. It is hoped that the rigorous study of perceptual learning will contribute to this momentum.

Acknowledgments

Thanks to the participants of my fall 1992 graduate seminar on "Plasticity and Cognition" for helpful discussions, especially colleagues Paul Bloom and William Ittelson. Thanks also to Doug Medin for helpful editorial comments, and to Elizabeth Niswander for dozens of hours of figure preparation. This work was supported by a grant from the National Science Foundation (BNS-8909825).

REFERENCES

Adams, J. A. (1987). Historical review and appraisal of research on the learning, retention, and transfer of human motor skills. *Psychological Bulletin, 101,* 41–74.

Allan, L. G., & Siegel, S. (1986). McCollough effects as conditioned responses: Reply to Skowbo. *Psychological Bulletin, 100,* 388–393.

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

Bedford, F. L. (1992a). *Constraints on perceptual learning: Objects and dimensions.* Submitted for publication. Manuscript available from author.

Bedford, F. L. (1992b). *Coordinating spatial maps: Rules of extrapolation.* Submitted for publication.

Bedford, F. L. (1993). Perceptual and cognitive spatial learning. *Journal of Experimental Psychology: Human Perception and Performance.*

Bedford, F. L., & Reinke, K. S. (in press). The McCollough effect: Dissociating retinal from spatial coordinates. *Perception and Psychophysics.*

Berwick, R. (1985). *The acquisition of syntactic knowledge.* Cambridge, MA: MIT Press.

Bloom, P. (1990). [Review of *Cognition and representation*] *Mind & Language, 5,* 166–173.

Brand, J. L., Holding, D. H., & Jones, P. D. (1987). Conditioning and blocking of the McCollough effect. *Perception and Psychophysics, 41,* 313–317.

Breitmeyer, B. G., & Cooper, L. A. (1972). Frequency specific color adaptation in the human visual system. *Perception and Psychophysics, 11,* 95–96.

Cassirer, E. (1944). The concept of group and the theory of perception. *Philosophy and Phenomenological Research, 5,* 1–36.

Chen, L. (1985). Topological structure in the perception of apparent motion. *Perception, 14,* 197–208.

Cheng, K., & Gallistel, C. R. (1984). Testing the geometric power of an animal's spatial representation. In H. Roitblat, T. G. Bever & H. Terrace (Eds.), *Animal cognition* (pp. 409–423). Hillsdale, NJ: Lawrence Erlbaum.

Chomsky, N. (1986). *Knowledge of language: Nature, origin, and use.* New York: Prager.

Cunningham, H. A. (1984). An Apple microcomputer-based laboratory for the study of visual–motor behavior. *Behavior Research Methods, Instruments, and Computers, 17,* 484–488.

Dodwell, P. C. & Humphrey, G. K. (1990). A functional theory of the McCollough effect. *Psychological Review, 97,* 78–89.

Ellis, S. R. (1976). Orientation constancy of the McCollough effect. *Perception and Psychophysics, 19,* 183–192.

Epstein, W. (1967). *Varieties of perceptual learning.* New York: McGraw-Hill.

Epstein, W. (1976). Recalibration by pairing: A process of perceptual learning. *Perception, 3,* 59–72.

Fodor, J. A. (1983). *Modularity of mind.* Cambridge, MA: MIT Press.

Foreit, K. G., & Ambler, B. A. (1978). Induction of the McCollough effect I: Figural variables. *Perception and Psychophysics, 24,* 295–230.

Gallistel, C. R. (1990). *The organization of learning.* Cambridge, MA: MIT Press.

Gibson, E. J. (1969). *Principles of perceptual learning and development.* New York: Appleton-Century-Crofts.

Graham, N. (1992). Breaking the visual stimulus into parts. *Current Directions, 1,* 55–61.

Harris, C. S. (1965). Perceptual adaptation to inverted, reversed, and displaced vision. *Psychological Review, 72,* 419–444.

Harris, C. S. (1980). Insight or out of sight? Two examples of perceptual plasticity in the human adult. In C. S. Harris (Ed.), *Visual coding and adaptability* (pp. 95–149). Hillsdale, NJ: Lawrence Erlbaum.

Hay, J. C. (1974). Motor-transformation learning. *Perception, 3,* 487–496.

Hay, J. C., & Pick, H. L., Jr. (1966). Gaze-contingent prism adaptation: Optical and motor factors. *Journal of Experimental Psychology, 72,* 640–648.

Held, R. (1965). Plasticity in sensory–motor systems. *Scientific American, 213,* 84–94.

Held, R. (1980). The rediscovery of adaptability in the visual system: Effects of intrinsic and extrinsic chromatic dispersion. In C. S. Harris (Ed.), *Visual coding and adaptability* (pp. 69–94). Hillsdale, NJ: Lawrence Erlbaum.

Held, R. & Durlach, N. (1991). Telepresence, time delay, and adaptation. In S. R. Ellis (Ed.), *Pictorial communication in virtual and real environments* (pp. 232–246). London: Taylor and Francis

Held, R. & Shattuck, S. R. (1971). Color and edge-sensitive channels in the human visual system: Tuning for orientation. *Science, 174,* 314–316.

Helson, H. (1964). *Adaptation-level theory: An experimental and systematic approach to behavior.* New York: Harper & Row.

Hepler, N. (1968). A motion-contingent aftereffect. *Science, 162,* 376–377.

Howard, I. P. (1982). *Human visual orientation.* Chichester, New York: John Wiley & Sons.

Humphrey, G. K., Dodwell, P. C., & Emerson, V. F. (1985). The roles of pattern orthogonality and color contrast in the generalization of pattern-contingent color aftereffects. *Perception and Psychophysics, 38,* 343–353.

Hurvich, L. M., & Jameson, D. (1957). An opponent-process theory of color vision. *Psychological Review, 64,* 384–404.

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology,* 175–219.

Keil, F. C. (1992). The origins of an autonomous biology. *The Minnesota Symposium on Child Psychology, 25.* Hillsdale, NJ: Lawrence Erlbaum.

Klein, F. (1957). *Vorlesungen uber hohere geometrie* [Lectures on higher geometry] (3rd ed.). New York: Chelsea. (Original work published 1893)

Koh, K., & Meyer, D. E. (1989). Induction of continuous stimulus-response relations. In G. Olsen & E. Smith (Eds.), *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society* (pp. 333–340). Hillsdale, NJ: Erlbaum.

Kohler, I. (1964). The formation and transformation of the perceptual world. *Psychological Issues, 3*(4), 1–173.

Levine, M. (1982). You-are-here maps: Psychological considerations. *Environment and Behavior, 14,* 221–237.

Lorenz, K. Z. (1981). *The foundations of ethology.* New York: Springer-Verlag.

Lovegrove, W. J., & Over, R. (1972). Color adaptation of spatial frequency detectors in the human visual system. *Science, 176,* 541–543.

Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science, 14,* 57–77.

Mayhew, J. E. W., & Anstis, S. M. (1972). Movement aftereffects contingent on color, intensity, and pattern. *Perception and Psychophysics, 12,* 77–85.

McCollough, C. (1965). Color adaptation of edge detectors in the human visual system. *Science, 149,* 1115–1116.

Modenov, & Parkhomenko (1965). *Geometric transformations* (Vols. 1 & 2). New York: Academic Press.

Murch, G. M. (1976). Classical conditioning of the McCollough effect: Temporal parameters. *Vision Research, 16,* 615–619.

Redding, G. M., & Wallace, B. (1988a). Adaptive mechanisms in perceptual–motor coordination: Components of prism adaptation. *Journal of Motor Behavior, 20,* 242–254.

Redding, G. M., & Wallace, B. (1988b). Components of prism adaptation in terminal and concurrent exposure: Organization of the eye–hand coordination loop. *Perception and Psychophysics, 44,* 59–68.

Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist, 43,* 151–160.

Rescorla, R. A., & Holland, P. C. (1976). Some behavioral approaches to the study of learning. In M. R. Rozenzweig and E. L. Bennett (Eds.), *Neural mechanisms of learning and memory.* Cambridge, MA: MIT Press.

Rozin, P. (1976). The evolution of intelligence and access to the cognitive unconscious. In J. A. Sprague & A. N. Epstein (Eds.), *Progress in psychobiology and physiological psychology: Vol. 6* (pp. 245–280). New York: Academic Press.

Rozin, P., & Schull, J. (1988). The adaptive-evolutionary point of view in experimental psychology. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Handbook of experimental psychology* (pp. 503–546). New York: Wiley-Interscience.

Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology, Learning, Memory and Cognition, 13,* 501–518.

Schacter, D. L. (1992). Understanding implicit memory: A cognitive neuroscience approach. *American Psychologist, 47,* 559–569.

Shepard, R. N. (1984). Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review, 91,* 417–447.

Shepard, R. N. (1987). Evolution of a mesh between principles of the mind and regularities of the world. In J. Dupre (Ed.), *The latest on the best: Essays on evolution and optimality.* Cambridge, MA: MIT Press.

Shepard, R. N. (1989). Initial representation of universal regularities: A challenge for connectionism. In L. Nadel, P. Culicover, L. A. Cooper, & R. M. Harnish (Eds.), *Neural connections, mental computations.* Cambridge, MA: MIT Press/Bradford Books.

Shepard, R. N. (1991). *Can natural selection yield universal principles of mind?* Paper presented at AAAS, Washington, DC, Feb. 1991.

Shepard, R. N. (1992). The three-dimensionality of color: An evolutionary accommodation to an enduring property of the world. In J. Burkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 495–536). New York: Oxford University Press.

Siegel, S., & Allan, L. G. (1985). Overshadowing and blocking of the orientation-contingent color aftereffects: Evidence for a conditioning mechanism. *Learning and Motivation, 16,* 125–138.

Siegel, S., & Allan, L. G. (1987). Contingency and the McCollough effect. *Perception & Psychophysics, 42,* 281–285.

Siegel, S., & Allan, L. G. (1992). Pairings in learning and perception: Pavlovian conditioning and contingent aftereffects. In D. Medin (Ed.), *Psychology of learning and motivation, 28,* 127–155.

Siegel, S., Allan, L. G., & Eissenberg, T. (1992). The associative basis of contingent color aftereffect. *Journal of Experimental Psychology: General, 121,* 79–94.

Skowbo, D. (1984). Are McCollough effects conditioned responses? *Psychological Bulletin, 46,* 215–226.

Skowbo, D., & Forster, T. (1983). Further evidence against a classical conditioning model of McCollough effects. *Perception and Psychophysics, 34,* 552–554.

Skowbo, D., Timney, B. N., Gentry, T. A., & Morant, R. B. (1975). McCollough effects: Experimental findings and theoretical accounts. *Psychological Bulletin, 82,* 497–510.

Sloane, M. E., Ost, J. W. P., Etheridge, D. B., & Henderlite, S. E. (1989). Overprediction and blocking in the McCollough aftereffect. *Perception and Psychophysics, 45,* 110–120.

Spelke, E. S. (1990). Principles of object perception. *Cognitive Science, 14,* 29–56.

Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review, 99,* 605–632.

Stromeyer, C. F. (1978). Form–color aftereffects in human vision. In R. Held, H. Leibowitz & H. L. Teuber (Eds.), *Handbook of sensory physiology, Vol. 8. Perception* (pp. 97–142). New York: Springer-Verlag.

Taub, E. (1968). Prism compensation as a learning phenomenon: A phylogenetic perspective. In R. Held, E. Taub, & S. J. Freedman (Eds.), *Neuropsychology of spatially oriented behavior* (pp. 77–106). Chicago, IL: Dorsey Press.

Taylor, J. G. (1962). *The behavioral basis of perception.* New Haven, CT: Yale University Press.

Wallach, H. (1968). Adaptation based on cue discrepancy. R. Held, E. Taub, & S. J. Freeman (Eds.), In *Neuropsychology of spatially oriented behavior* (pp. 235–259). Illinois: Dorsey Press.

Warren, R. M. (1985). Criterion shift rule. *Psychological Review, 92,* 574–584.

Warren, W. H. (1977). Visual information for object identity in apparent movement. *Perception and Psychophysics, 25,* 205–208.

Welch, R. B. (1971). Discriminative conditioning of prism adaptation. *Perception and Psychophysics, 10,* 90–92.

Welch, R. B. (1972). The effect of experienced limb identity upon adaptation to simulated displacement of the visual field. *Perception and Psychophysics, 12,* 453–456.

Welch, R. B. (1978). *Perceptual modification.* New York: Academic Press.

Welch, R. B. (1986). Adaptation of space perception. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance* (Vol. 1) (pp. 24–45). New York: Wiley-Interscience.

Welch, R. B., & Warren, D. H. (1980). Immediate perceptual response to intersensory discrepancy. *Psychological Bulletin, 88,* 638–667.

Westbrook, R. F., & Harrison, W. (1984). Associative blocking of the McCollough effect. *Quarterly Journal of Experimental Psychology, 36A,* 309–318.

White, K. D., and Riggs, L. A. (1974). Angle-contingent color aftereffects. *Vision Research, 14,* 1147–1154.