

Perceptual and Cognitive Spatial Learning

Felice L. Bedford

Subjects were taught novel mappings between visual space and motor space with either a variant on a prism adaptation paradigm (Experiments 1 and 2) or a nonperceptual cognitive task (Experiments 3 and 4). First, discrimination training specified that 1 visual location required a new pointing response but another location did not. This led to unusual generalization unlike typical generalization decrement. Second, training at 9 locations specified that 1 location required a new response but that the remaining 8 did not. This simple isolation mapping was unlearnable and instead a flat function fit through all of space. In contrast, for the cognitive paradigm, not only was isolation of one region of space easily learned, it was the preferred pattern of generalization. Implications for perceptual learning, as well as the qualitative distinctions between perceptual and cognitive learning, are discussed.

Many judgments in perception can be made with more than one system. For instance, consider trying to locate an object in space. Is it to your right? Your left? Straight ahead? One can decide visually by determining the position of the object on the retina and the position of the eyes in the head. Alternatively, one can do so without vision by groping for the object with the hand and relying on information from the joints. If the object is noisy, then the auditory system provides yet another alternative. The ability to use several independent systems is not limited to spatial location; size and orientation, for instance, can also be determined both with the eyes and with the hand. The fact that many, if not all, perceptual parameters are multiply determined implies that a fundamental question in perception is how the information from the different sources gets coordinated or aligned (cf. Wallach, 1976; Wallach & Karsh, 1963a).

Specifically, the different systems must be kept in perfect alignment to perceive a coherent world. Imagine the confusion if your eyes indicated that a desired object was to your right, but your ears localized it to your left. Where is the object really? Where should you reach to obtain it? I argue that the reason this situation would be disconcerting is because in our physical world an object can only be in one place at any one time. Otherwise, detecting an object in multiple locations seems unremarkable. Moreover, it is likely that such a pervasive constraint on our world has been internalized by our perceptual systems (cf. Shepard, 1984, 1987,

1991, 1992). Therefore, if our sensory systems detect that an object is located in two places at once, nothing will be learned about the state of the world. Instead, an internal malfunction will be inferred and subsequently corrected. This set of events can be demonstrated artificially through the now classic phenomenon of *prism adaptation*, in which a wedge prism causes the visual image to shift laterally with respect to the other sensory systems. Outside of the laboratory, discrepancy between sensory systems may occur during growth (e.g., Held, 1965) if parts of the body subserving those systems (e.g., arm length, interocular distance) change size and shape by different amounts or at different rates. In addition, in adults, fine tuning may be continually necessary because complex systems can drift with respect to one another without an active process of coupling (Howard, 1982). The alignment necessary for our usual internal coherence then cannot be completely prewired. Learning, as broadly defined, must be involved (see Bedford, in press).

What are the rules that characterize this general learning process? One salient yet neglected feature is that entire dimensions of stimuli must be connected (Bedford, 1989). For instance, consider an array of positions from far left to far right as they are localized visually and the same positions as they are localized with the arm, or proprioceptively. Because an object can appear anywhere along this continuum, you need to know exactly how to direct your hand for every possible position. Consequently, a whole mapping between a visual and a motor dimension is required. Is learning a mapping between dimensions of stimuli fundamentally different from learning to connect a single pair of stimuli?¹

Felice L. Bedford, Department of Psychology, University of Arizona.

This research was supported by National Science Foundation Grant BNS-8909825. Experiment 2 was presented at "Orientation in Space," the 17th Center for Visual Sciences Symposium, June 21-23, 1990, University of Rochester, New York.

I thank William Ittelson and reviewers Gordon Redding and Robert Welch for helpful comments, and Karen Reinke and Chris Bockisch for assisting with data collection.

Correspondence concerning this article should be addressed to Felice L. Bedford, Department of Psychology, University of Arizona, Tucson, Arizona 85721. Electronic mail may be sent to bedford@rvax.ccit.arizona.edu.

¹ A more precise statement is "detecting a mapping." The internal constraint implying that objects cannot be in two places at once may never be suspended. Consequently, any new mapping that indicates that objects are localized in one place visually but localized in a different place proprioceptively is not learned but is instead detected, and the information is filtered down to the visual and proprioceptive subsystems. It is within these subsystems that the change occurs (e.g., Harris, 1965) to restore alignment. For convenience, the phrase *learning a mapping* will be used.

Logically, a mapping can be decomposed into a list of connections between individual visual positions and accompanying motor–proprioceptive positions ($v_1, p_1; v_2, p_2; \dots, v_n, p_n$, where v_n = a position along the visual dimension, and p_n = a position along the proprioceptive dimension). Does the perceptual system do this as well? Whether learning a mapping is an instance of straightforward associative learning is important for two general theoretical issues, one in perception and one in learning. In perception, understanding the rules by which dimensions are coordinated has been an issue at least since Berkeley (1709/1910) discussed the acquisition of depth perception. In learning, the question of whether there are a few general learning mechanisms or many specific ones has recently resurfaced (e.g., Rozin & Schull, 1988).

How useful the framework of elementary associative learning is for understanding perceptual mappings is not well known. In learning, research has been dominated by study of the formation of a single associative connection. Traditional Pavlovian and instrumental learning processes assume that the fundamental building block of learning is a single connection between a stimulus and response or between a stimulus and another stimulus. Yet it has been noted that connecting a large number of stimuli to a large number of responses may not necessarily be predictable from knowing all the consequences of associating just a single stimulus and response (e.g., Carroll, 1963). In perception, the classic experimental strategy of studying spatial mappings goes to the opposite extreme of pairing large continuous regions of stimuli and responses. In prism adaptation, vision is artificially displaced with respect to touch by looking through a prism. A subject is exposed to the new visual–motor arrangement by looking at his or her hand moving back and forth through the prism, for instance (for reviews, see Howard, 1982; Welch, 1978, 1986). Because there are many possible visual–proprioceptive pairs of locations defined with this paradigm, it has been difficult to analyze a mapping and investigate its component parts. Consequently, the apparent relationship between perceptual adaptation and elementary learning processes has remained largely untested, though it has been the subject of continuing interest (see Bedford, 1989; Epstein, 1975; Taub, 1968; Wallach, 1984; Welch, 1978).

Recently, a new procedure was developed to decompose the spatial dimensions experimentally (Bedford, 1989). Training occurs in the dark with a small red light-emitting diode (LED) worn on the finger, which is initially unilluminated. A different LED, one located in space, is illuminated, and a subject is told that only while he or she points accurately to the light in space will the light on his or her finger turn on. Consider that with undistorted vision, this is an easy task: With our normal coordination it is easy to direct the arm, even without visual guidance, to an object in space, which would then cause the LED to turn on (Figure 1a). But the subject looks through a prism, making the target appear to the side. Therefore, the subject points to the side of the object's true location, failing to illuminate the finger LED (Figure 1b). After a subject is informed that it is "disorienting being in the dark," he or she moves his or her arm around and quickly stumbles on the correct position, which causes

the finger LED to light up. When this happens, the subject localizes the finger light proprioceptively in one location but sees it (also through the prism) in a different location (Figure 1c). When the finger LED is illuminated then, training is provided with a single visual–proprioceptive pair of locations. A variable prism under computer control allows any offset to be chosen, which creates a flexible paradigm whereby individual regions can be controlled separately by the experimenter.

Experiments conducted with this procedure suggest that a mapping between spatial dimensions is not, in fact, represented as a list of connections between individual locations. In one experiment, subjects were trained with multiple trials of only a single new pair of locations (e.g., a visual position straight ahead that must be pointed to at a location a few inches to the left). If one pair of locations is the correct unit of analysis, then each pair is relatively independent of the others, and training at one location would be expected to have minimal effects elsewhere. There should be a declining impact as we move away from the trained location, analogous to generalization decrements found in other domains when one pair of stimuli are associated (e.g., Macintosh, 1974). Instead, training at one location produced a rigid shift throughout the entire dimension such that pointing shifted everywhere by the same amount. Instead of choosing to isolate one region, the entire dimension was "picked up" and treated as the unit instead (Bedford, 1989).

Does the preference for a global realignment imply that one small region can never be isolated? In many domains, constraints are rarely all or none. Instead, a hierarchy of preferences can be revealed. As part of a larger goal to fully elucidate the constraints on learning new mappings between perceptual dimensions, the purpose of the present study was to determine whether there are any conditions under which a pair of small spatial locations would be treated as an independent unit separately manipulable from adjacent regions. Experiments 1 and 2 provided training conditions increasingly likely to produce such isolation. A secondary purpose was to contrast this process with more familiar learning processes in which individual stimuli are readily treated separately. Would it be possible to find qualitative differences between perceptual and cognitive learning processes? Experiments 3 and 4 used a cognitive learning paradigm while repeating the mappings of the first two experiments.

Experiment 1

Training a single visual location with a new pointing response leads to complete generalization of that response to all visual locations. Complete generalization, however, may not occur if training at a second location were added such that the pointing response required at that location was simply the normal response used in everyday life. For instance, imagine that when you see an object straight ahead you must reach 5 in. (12.7 cm) to the right to secure it, but when the object is to your left you reach exactly where you always have, directly at the seen location. It would now be illogical to infer that you should reach 5 in. (12.7 cm) to the right regardless of where you see the object. Though complete generalization




Figure 1. Finger LED procedure shown for successful pointing with undistorted vision (Panel a), initial attempt with prism-displaced vision (Panel b), and successful pointing with prism-displaced vision (Panel c). (The sunbursts show the seen positions of the light-emitting diodes. The filled circles show the actual positions, if different. *P* = felt location of finger [proprioceptive]. *V* = seen location of finger [visual].)

may be a rule on the top of a preference hierarchy, training that disproves the rule may uncover less preferred options. For the scenario described, one sensible conclusion would be that only one region of space has been changed. Interspersing normal experience at a second location then might encourage the isolation of one small region of space, if that option is available. The two visual locations used were straight ahead and 15° to the left. For the straight-ahead target, the required response was 10° to the right, and for the left target the pointing response was offset 0° , that is, not at all. To assess the effects of training, pointing was tested on 10 different positions along the left-right continua, including the 2 trained positions and 8 new ones. If one region is isolated, there should be a generalization gradient surrounding the straight-ahead position. To control for any effects of visual asymmetry during training (Ebenholtz, 1974; Paap & Ebenholtz, 1976), as well as training in general, another group of subjects was trained on the same two visual positions except that the offsets for both were 0° .

Method

Subjects. Sixteen undergraduate students that either had uncorrected vision or wore contact lenses received \$4.50 for partic-

ipating in the study. Ten of those subjects were women, and all were right handed.

Apparatus. The apparatus is described and pictured in more detail elsewhere (Bedford, 1989). The subject faced a panel of LEDs located every 2.5° from a distance of 111.8 cm. Pointing responses were measured with a rod that was fixed at a point located under a subject's midline. Angular position of the pointer was measured with a potentiometer. A cuff containing a red LED was worn on a subject's right index finger and was attached with a thin wire to the pointer. Subjects looked through a variable prism mounted in the right eyepiece of a stationary goggle-type frame, which was attached to a stepping motor. The distortion of the prism, the illumination of the LEDs in both the panel and on the finger, and measurement of the pointing responses were all controlled with an IBM XT computer. For all of the experiments, an adjustable chin rest was used for the majority of the subjects. The chinrest in combination with the stationary goggle frame kept a subject's head from moving. For 5 of the subjects, a bite bar was substituted when the chinrest needed repair.

Procedure. Subjects were assigned randomly to either the experimental or the control group. Seat height and chinrest were adjusted such that each subject could just barely see the "centering lights," one flickering LED at each of the two extreme left and right ends of the prism's viewing field. The centering lights were extinguished and were used again following all breaks. All room lights

were extinguished and remained off for the duration of the experiment. There were four major parts: normal pretraining, pretest, training, and posttest.

Pretraining of normal coordination (see Wallach & Karsh, 1963b) within the apparatus was used to minimize any initial pointing biases and to provide practice. With vision undistorted, subjects pointed to LEDs in 10 different positions ranging from 20° to the left to 25° to the right: -20°, -15°, -10°, -5°, 0°, 5°, 10°, 15°, 20°, and 25°. The targets were presented one at a time in a random order with two repetitions for each. Visual feedback of the finger while pointing was provided by turning on the finger LED at the start of the phase and leaving it on for the entire phase. Because of the ambiguity of the instruction "point to a target" (as opposed to grabbing it), subjects were instructed more specifically to point such that their right eye, the tip of the finger, and the light in space all lined up. Each subject pushed a button when satisfied and was required to swing his or her arm back and forth between trials while still using the pointer. This was done to maximize the likelihood that each trial was an independent attempt to point at a target. The intertrial interval was 2.4 s.

The pretest phase measured subjects' accuracy at pointing to targets without any visual feedback of the hand. The procedure was the same as normal pretraining except that the finger LED was turned off before starting and remained off. A subject was instructed to point such that: "It feels like your right eye, the tip of your finger, and the light in space are all lined up, even though you can't see the tip of your finger." They were also informed that it would probably be more difficult than the last part and to do the best they could. Two subjects were eliminated and replaced after this phase for having mean pointing biases of $\pm 5^\circ$ or more.

The training phase provided training with either the experimental or the control mapping. As in the prior two phases, LEDs in space were turned on one at a time. In this part, the finger LED was unilluminated at the start of each trial. The finger LED would be turned on when a subject successfully pointed to the target's true location. The subject's task was to light the finger LED and keep it illuminated as best he or she could for the duration of each trial, while the target LED remained on independent of a subject's performance. The finger LED would turn off if the pointer strayed by 1/4° on either side. After each trial, both the finger LED and the LED in space were turned off nearly simultaneously. The setting of the prism depended on which target was illuminated and which group the subject was assigned to. In both groups only two apparent visual locations were used, -15° and 0°. In the experimental group, the prism was set to 0° offset (undistorted) for the -15° target. For the other target, the prism was set to a 10° offset such that the visual image was shifted 10° to the left. To obtain a target that appeared visually at 0°, an LED actually located at 10° was used. (Note that to illuminate the finger LED for this target, a subject must point 10° to the right of the apparent 0° visual location.) In the control group, the offset for both targets was 0°. The setting of the prism was changed between trials. To go between a 0° offset and a 10° offset required 232 steps of the stepping motor, which took 0.9 s. To ensure equal information for all trials and groups, the stepping motor was always moved 232 steps between trials, returning back to the same position if necessary. The total intertrial interval was 2 s.

There was an initial practice block in which a subject had to meet a criterion of three successful finger illuminations for each of the two positions. The experimenter remained in the room with the subject to provide instructions. During the practice phase, each trial lasted 8.2 s. Following practice, subjects received an additional 96 training trials, 48 at each position. The training was divided equally into 2 blocks of trials with a short break (in darkness) between blocks. Trials were randomized within each block with the restriction that no more than 3 trials of the same type occurred in suc-

cession. Each trial lasted 6.6 s. Pilot work suggested that after the practice phase, subjects needed less time to execute a correct response. Each subject was required to swing his or her arm back and forth after each trial.

The posttest phase assessed how training affected pointing to targets. The procedure used was identical to the pretest phase. Before beginning, however, subjects were told: "Don't worry about anything you may have learned in the last parts." Instructions for pointing were repeated: "Just point so that it feels like your right eye, the tip of your finger, and the light in space are all lined up." The instructions were included to ensure that any strategy a subject used in training would not persist in the testing phase.

Results

The pretest phase found comparable baseline pointing for the experimental and control groups, which received identical treatment up to this point. Data for perfectly accurate pointing would fall along the line $P = 1V + 0$, which reflects normal coordination. The best fit lines for the experimental and control groups are $P = 1.04V - 0.86$ ($SE = 0.03, 0.76$) and $P = 1.06V + 0.36$ ($SE = 0.03, 0.62$). The slopes with values greater than 1 reflect a tendency in both groups to begin with a pointing range slightly expanded with respect to the visual range. Because of the similarity of the two groups and no major deviations from normal coordination, all subsequent analyses were performed on the change in pointing from pretest to posttest. The change in pointing as a function of target position can be seen for both groups in Figure 2. The experimental and control groups differ, suggesting that the two-pair experimental training had an effect above and beyond any effects due to visual asymmetry or general training alone. This is confirmed by a significant group main effect, $F(1, 14) = 5.9, p < .05$, in an analysis of variance (ANOVA) performed on changes in pointing with group (experimental, control) and target (-20 through 25) as variables. A marginally significant effect of target, $F(9, 126) = 1.8, p < .10$, and a Target \times Group interaction, $F(9, 12) = 1.8, p < .10$, suggest that pointing differed for different targets, with a different pattern for each group.

The trends were analyzed further for each group. In the experimental group, Figure 2 suggests that there was no generalization gradient surrounding the straight-ahead target. If there was, the size of the pointing response should fall off on both sides of 0, including the region on the right from straight ahead (0°) to 25°. Instead, pointing is on average flat in that region with a slope of 0.003, not significantly different from 0 ($t < 1$). Over the entire range, from -20° through 25°, there were significant linear components, $F(1, 63) = 11.8, p < .01$, and quadratic components, $F(1, 63) = 5.6, p < .025$, and no significant cubic or higher order components ($F_s < 1$). (For details of this analysis, see Bedford, 1989.) The quadratic component reflects one "bend" in the data, or two distinct parts of the curve. These results suggest that there may have been a linear increase in the size of the pointing response between the two trained positions that then levels off in the extrapolative region, a pattern reported elsewhere (Bedford, 1989, in press).

In the control group, there were also changes in pointing. The up and down pattern as a function of target position seen

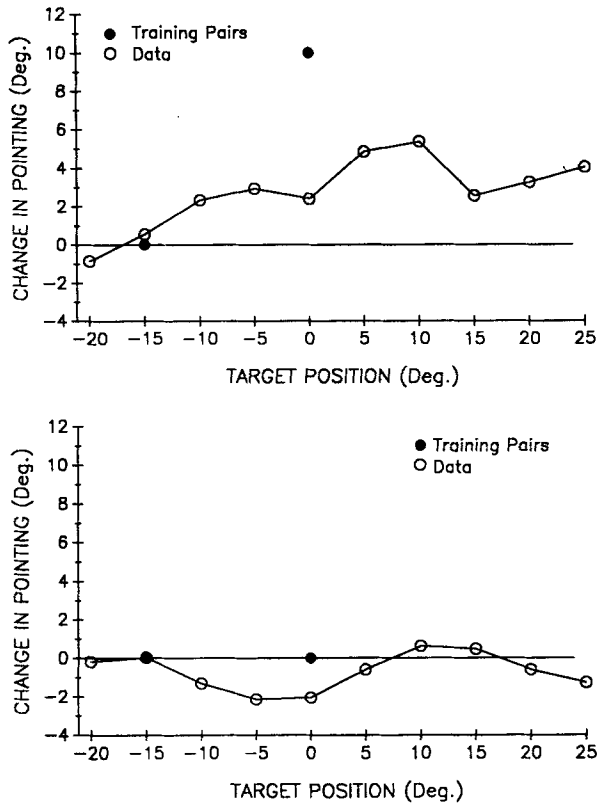


Figure 2. Change in pointing following two-pair training for the experimental group (top) and the control group (bottom) in Experiment 1.

in Figure 2 manifests itself as a significant cubic component, $F(1, 63) = 4.3, p < .05$. A constant mean leftward shift of 0.7° was not significant ($t < 1$), and no other trend was significant (all $F_s < 1$). Other reported effects of visual asymmetry (Ebenholtz, 1974; Paap & Ebenholtz, 1976) are of constant shifts, but the data are reported for only one test position or collapsed across all target positions. Whether the current pattern is due to visual asymmetry or some other aspect of training is not known. Nonetheless, the pattern of results in the experimental group is visible even with the data from the control group subtracted (Figure 3).

Discussion

Subjects received training that one region of space remained normal along with training that another region was shifted. This additional clue, however, did not cause the one shifted region to be isolated. One way to view the present experiment is as a discrimination experiment: One stimulus requires a new response, and another stimulus does not. Although discrimination training in other domains produces particularly sharp generalization decrements (e.g., Sutherland & Macintosh, 1971, p. 220), here the size of the new response did not diminish as the distance from the shifted training position increased. Instead, the inference made about untrained regions of space is arguably more complex: For

positions to the left of straight ahead, the change in pointing did decrease gradually, but to the right of straight ahead, pointing remained shifted by the same amount. Rather than isolating one region, the two-pair training had more global effects throughout the entire dimension. The data are consistent with other two-pair training manipulations in which interpolation between the two trained positions was found to be a sloping line (Bedford, 1989), and extrapolation beyond them appeared to be flat (Bedford, in press). In the present experiment, positions to the left of straight ahead lie within the interpolative region, and positions to the right are in the extrapolative region. Apparently, the rules apply even for this special case in which one of the two required responses is unchanged from normal, suggesting that the rules are highly preferred and applicable to a wide variety of circumstances. For the present purposes, the data suggest that one small region of space was not treated independently, though intuitively it would have been a simple and sensible approach.

Experiment 2

Note that although one region was not isolated, the solution that was chosen was not incorrect. Because there are only two training pairs, any of the infinite number of functions that goes through those two points would accommodate the experimental conditions equally well. Consequently, if dividing up the dimension into small independent regions were an option low on a preference hierarchy, then this ability may not be revealed when other options can be used instead. The next experiment provided further encouragement by using a mapping in which just about any interpretation other than "only one region has changed" would be a failure to learn correctly. To accomplish this, the number of training pairs was increased from two to nine. Subjects were again trained on a single shifted location at straight ahead, but now instead of just one explicitly trained normal region, eight more normally trained regions were added. A second group of subjects was trained on the reverse mapping: "All but one region has changed." That is, now eight positions were trained with a shift, whereas the straight-ahead position required normal pointing. Although training with only a couple of pairs is




Figure 3. Change in pointing after the control group is subtracted from the experimental group in Experiment 1.

ambiguous and therefore can reveal preferences, training that unambiguously specifies a mapping can be used to search for absolute limitations on what can be learned. Although logically an infinite number of functions can still fit through nine pairs, psychologically I considered this training to constitute a complete mapping of the enclosed space.

Can mappings between the visual and proprioceptive spatial dimensions be learned when they require that one region be singled out? And if not, what will happen to pointing? The answers to these questions are unknown, in part because previous research on transformations has been limited to the mappings imposed by prisms and other optical devices. The present finger LED paradigm has the flexibility to allow many rearrangements of interest to be investigated.

Method

Subjects. Sixteen students (6 women, 10 men) enrolled in an introductory psychology course received course credit for participating. All subjects were right handed and had either uncorrected vision or wore contact lenses.

Procedure. The procedure was similar to the procedure of Experiment 1, with the following differences. During the training phase, both groups received training on visual positions at -20° , -15° , -10° , -5° , 0° , 5° , 10° , 15° , and 20° . There were again 96 trials; in this experiment, half of those were at 0° , and the other half were distributed equally among the other eight positions. For the "only one region is changed" mapping (Group 1), the 0° position was trained with a 10° rightward shift, and training everywhere else was with a 0° shift. For the "all but one region is changed" mapping (Group 2), the 0° position was trained with 0° shift, and training everywhere else was with a 10° rightward shift. That is, in Group 1, subjects must point to the right for the straight-ahead target and normally everywhere else, and in Group 2, subjects must point normally for straight-ahead targets and to the right everywhere else. The criterion for the practice block was successful illumination of the finger LED three times for the 0° position and three times total from any of the other positions. The only other difference from Experiment 1 was that the normal pretraining phase was not used.

Results and Discussion

It is clear from Figure 4 that these mappings were not well learned. Statistically, the two groups did not differ: In an ANOVA with group and target as variables, neither group, target, nor their interaction was significant (all $F_s < 1$). Had the mappings been correctly learned, there should have been a Group \times Target interaction. Further trend analysis on each group separately indicated a completely flat function in Group 1: No component (linear, quadratic, cubic, and higher) was found significant (all $F_s < 1$). In Group 2, there was a small significant quadratic component, $F(1, 63) = 8.3$, $p < .01$, suggesting that the mapping "all but one region has changed" was learned slightly better than the mapping "only one region has changed."

Note that although the mappings were not learned very well, behavior did change. The primarily flat functions were not at 0° but at a mean of 4.7° and 3.7° for Groups 1 and 2, respectively. That is, pointing shifted primarily rigidly by an amount approximately halfway between the two imposed




Figure 4. Change in pointing for the "only one region has changed" mapping (top) and the "all but one region has changed" mapping (bottom) in Experiment 2.

distortions. First, this suggests that all of the training did get processed. For instance, in Group 1, had there not been any change in pointing at all, it could have been because the only location in which there was a shift was simply ignored. But more important, the change implies that when the system is confronted with a mapping it finds difficult or unnatural, it does not simply shut down and do nothing. It accommodates behavior as best it can, subject to its own internal rules and structure. Consequently, the deviation between an imposed mapping and actual behavior can be used to reveal the rules and constraints on this learning system, as had been hoped (Bedford, 1989; Hay, Langdon, & Pick, 1971). The behavior that did occur provides converging evidence with my previous work on "incomplete" mappings (Bedford, 1989), which found a strong preference for linear functions. Here, the best fit linear function is a change in intercept only, that is, a rigid shift, which may have been imposed on the conditions of training.

Performance for the "all but one region has changed" mapping was better than for the "only one region has changed" mapping. If anything, however, it is the latter situation that seems more straightforward and simple. Most of the intuitions on how one would intentionally solve these problems have proved inapplicable. Although ultimately the differences in learnability between the two mappings may uncover further rules, the reason for the difference is not known at

present. In addition, it is possible that with further training both mappings would be learned better, an argument that can always be made in a learning experiment. The point is not that these distortions can never be accommodated but rather the difficulty they appear to present. A mapping in which one small region was singled out, especially a mapping in which only one new motor response was required, would be extremely easy to learn if the correct building block of learning was a small region of space. For the present purposes, then, this study suggests that a single stimulus-response connection does not behave as an independent unit. As well as preferring not to isolate one small 5° region of space, this perceptual learning system has great difficulty doing so even when the situation demands it.

There is a practical problem, however. The flat function can also be due to a low-level failure to discriminate visually between adjacent target positions. If the straight-ahead position was often mistaken with its neighbors on either side, then the training would be experienced as rather inconsistent. It might seem from the subjects' perspective as if the same position in space sometimes was shifted and sometimes not. Under such inconsistent and confusing training, an averaging strategy could be used. Similarly, there could be an uninteresting low-level motor restriction that prevents the correct arm movements from occurring even if the mapping had been correctly represented. The next experiment attempted to rule out these alternatives.

Experiment 3

A procedure was needed to show that the visual locations used in Experiment 2 could be discriminated and that there was no trouble producing the motor responses for those positions. Choice of a procedure was also motivated by surprise that such an intuitively simple mapping was not learned. In the previous experiment, subjects were never told to try to learn anything. Their task was to point to targets as accurately as they could, as well as to light up the finger LED, while out of awareness perceptual learning occurred. Imagine instead an intentional learning situation in which multiple stimuli were to be matched to multiple responses. If only one of these stimuli required a new response, the task would seem extremely easy, almost boring. It should not take long to figure out what was required, nor would it burden memory to remember. To test these intuitions, as well as to test visual and motor discrimination, a cognitive procedure was developed that required subjects to learn connections intentionally but did not involve perceptual learning.

To accomplish this, vision was never distorted with a prism. Instead, the same mappings from Experiment 2 were simulated with computer software. The difference between this procedure and the perceptual learning procedure is best illustrated by an example: Consider a light in space that appears 10° to the right of straight ahead but must be pointed to straight ahead to illuminate the finger LED. In both procedures, the target light appears at 10° , and the hand must be at 0° to be successful. In the perception procedure, when this occurs the hand looks as if it is also at 10° and looks as if it is pointing directly to the target (refer again to Figure 1c).




Figure 5. Depiction of successful performance in the cognitive procedure. (The sunbursts show the seen positions of the light-emitting diodes, which is the same as the actual positions. *P* = felt location of finger [proprioceptive]. *V* = seen location of finger [visual]. Compare with Figure 1c.)

This is because the target's actual position is at 0° , and both it and the hand are seen through the prism. In the cognitive procedure, however, in which no prism is involved, the hand simply looks where it actually is, at 0° , and furthermore looks as if it is pointing to the side of the target (Figure 5). Subjects were told from the beginning that this was a learning task; they were told that they had to try to figure out where to put their arm for each different light and that they would be tested on it afterward.

Critically, in the new procedure there is no discrepancy between where the finger feels it is and where it appears to be. The seen and felt positions are always identical. In the previous example, when the hand is at 0° , both the visual and proprioceptive systems agree that the hand is at 0° . Therefore, there is no internal inconsistency between sensory systems that must be corrected.² Without the need for correction, the motivation for a perceptual learning process is absent. Instead, this is an abstract cognitive task. The task does, however, use the same visual stimuli and motor responses required in Experiment 2. Failure to learn the mappings in this experiment would raise additional questions, and success would accomplish two things: It would demonstrate that the failure in Experiment 2 was not due to low-level limitations, and moreover it would serve to contrast the perceptual learning process to different learning mechanisms.

Method

Subjects. Sixteen undergraduates (7 women, 9 men) received either course credit or \$4.50 for participating. All subjects were right handed and had normal or contact-corrected vision.

² The discrepancy between the sensory systems is critical, not the discrepancy between a sensory system and where the object is. For instance, if the hand were really at 10° but both the visual and proprioceptive systems detected it at 0° , there still would not be any discrepancy—unless a third sensory system, such as the auditory system, did detect it at 10° .

Procedure. Each subject was informed at the beginning of the experiment that this was going to be a learning experiment and was instructed to figure out and remember where to put his or her hand for each light position. Subjects were also told that they would be tested in two ways: reproducing the arm positions and verbally describing what they had learned. Subjects were assigned to either the "only one region has changed" or the "all but one region has changed" group. The mappings were created without the prism, as described earlier. The experiment proceeded immediately to the practice training phase, skipping the initial testing phase. It was thought that starting with a pointing task that had nothing to do with learning might misdirect the subjects' attention and interfere with the purpose of the experiment. During the training phase, half as many training trials were used as in Experiment 2. If subjects did find it easy to learn, too many extra trials might cause them to lose interest. After training, subjects were told that they were being tested and were instructed to use what they learned. For each target presented, they were to put their right arms in the position they thought would have illuminated the finger LED (Posttest 1). Following that phase, there was a second test (Posttest 2) in which subjects were told that this time they should just point directly to the target. The difference between Posttest 1 and Posttest 2 served in this experiment as the measure of the change in pointing strategy. After the experiment, subjects were asked to describe what they learned.

Results and Discussion

Figure 6 shows the change in pointing strategy, in which it is apparent that both mappings were learned fairly well. An ANOVA with group and target as variables confirms this: There was a significant Group \times Target interaction, $F(9, 126) = 4.8, p < .001$, which reflects the different pattern of pointing as a function of target position for the two groups. The target main effect was not significant, $F(9, 126) = 1.4, p > .1$, which is expected because the opposite target patterns of the two groups cancel when collapsed across group. The group main effect was significant, $F(1, 14) = 5.3, p < .05$, suggesting that the mean change for Group 2 was greater than for Group 1. Note that the pointing patterns were due to an intentional strategy and not any perceptual change. Although data on where subjects pointed before training were not collected, any perceptual effects that might change a subject's pointing have been removed by subtracting the second posttest, in which subjects were instructed to point directly at a target. Fifteen of 16 subjects were able to describe the major features of the mapping correctly; comments included the following. "One or a couple in center—move hand slightly to right; for others, directly under" (Group 1). "Line it up if on either side; if straight ahead move a little to the right of the light by 10° or 20°" (Group 1). "All but center, aim to right; aim straight at center" (Group 2). "Every light except one was to the right of light. Point straight at other one which was a little to the right" (Group 2). The 1 subject who was not able to describe the mapping correctly was from Group 1 and said that she learned to point straight at all of the lights on the left and point to the right for all of the lights on the right.

Although not perfect, the mappings were learned fairly well. This implies that the failure to learn the mappings in




Figure 6. Change in pointing strategy for the "only one region has changed" mapping (top) and the "all but one region has changed" mapping (bottom) using the cognitive procedure in Experiment 3.

Experiment 2 was not due entirely to an inability to make visual discriminations or motor movements. Had that been the only factor responsible for the results in Experiment 2, the data from that experiment would be expected to be at least as accurate as these. The difficulty must be due to the rules of the perceptual learning system itself. On the other hand, the mapping was found relatively easy to learn when presented as a problem-solving task, both verifying intuitions and contrasting the perceptual learning rules with cognitive abilities.

Experiment 4

The two learning systems did not handle the isolation mappings with equal ease. Perhaps, though, this is simply a parametric difference in rate of learning. Are the rules different enough to manifest different preferences as well? The two-pair training conditions of Experiment 1 were repeated, with the cognitive procedure replacing the perceptual learning procedure. Would the pattern of generalization to untrained positions be like Experiment 1? Or would this system, when given the option, choose instead to isolate one region?

Method

Subjects. Eight undergraduates (6 women, 2 men) with normal or contact-corrected vision received course credit for participating. None of the subjects participated in any of the previous experiments.

Procedure. Only the experimental group from Experiment 1 was used. During training, the straight-ahead position is shifted, and the -15° position is unshifted. As in Experiment 3, subjects were instructed to learn where to put their right hands for each light position, and the computer simulation procedure without the prism was used. In addition, Posttest 1 tested what subjects had consciously learned, and Posttest 2 measured pointing directly to the targets. All other aspects of the procedure were identical to Experiment 1, including normal pretraining, the pretest phase, and number of training trials.

Results and Discussion

The change in pointing strategy from pretest to Posttest 1 is shown in Figure 7 (top). It is clear that there is a generalization decrement, unlike the data from Experiment 1. The size of the pointing response falls off on both sides of the straight-ahead position. On the right side, from 0 to 25° , there is a negative slope of -0.38 , $t(7) = 2.19$, $p < .05$, unlike Experiment 1 in which it was flat in that region. The gen-

eralization gradient is also reflected by a quadratic component, $F(1, 63) = 13.3$, $p < .01$, in an analysis of trend across all positions (no other component significant). To ensure that the generalization gradient was due to intentional learning and not any unintentional perceptual changes in the seen or felt locations of objects, the analyses were repeated for data with Posttest 2 subtracted from Posttest 1. In the Posttest 2 phase, subjects were asked simply to point directly at the targets. There were no major differences in the analyses. Training did produce small perceptual-motor changes in pointing directly to targets, which were not unlike the changes in the control group of Experiment 1, which also presented asymmetric targets and no perceptual discrepancy. When these effects are subtracted, the gradient appears even more dramatic (Figure 7, bottom).

Subjects appear to have primarily singled out and isolated one position: The farther the distance of a novel position from the trained position, the less impact the trained response had, that is, a generalization gradient. These results are typical of the pattern of generalization following the association of a single stimulus and response and two-pair (discrimination) training in other domains. The data suggest that when cognitive inference is used, a single pair of locations behaves as an independent unit, unlike the perceptual learning system.

General Discussion

Can one small region of space be broken off and manipulated separately from other locations? The data suggest that for perceptual learning, this decomposition is not selected when it is a sensible option (Experiment 1) and, moreover, cannot be learned when explicitly trained (Experiment 2). In contrast, when general cognitive abilities are used, decomposition is not only easily learned (Experiment 3) but preferred (Experiment 4).

Dimension Learning

A previous study (Bedford, 1989) left open the possibility that under more favorable training conditions, adjacent regions could be manipulated separately for perceptual learning. The present study finds no evidence that such a mapping is learnable or that in general a mapping between spatial dimensions is ever represented as a list of discrete visual locations and corresponding motor locations. This finding was not a forgone conclusion. It is impossible to determine in principle which mappings are learnable or whether mappings can be reduced to a list of independent connections. This is true even if one considers what adaptation is for in the real world or what mappings will be encountered in the real world.

Howard (1982) summarized four possible reasons for the existence of the type of plasticity reflected by perceptual adaptation: recalibration for growth of the body, compensation for injury, correction of spontaneous drift between the senses, and adjustment to varied environments. For the last, he gave the example of a Japanese pearl diver who has a different stimulus world than that of a Greenland Eskimo.




Figure 7. Change in pointing strategy following two-pair training using the cognitive procedure in Experiment 4. (The top panel shows pretest subtracted from Posttest 1, and the bottom panel shows Posttest 1 subtracted from Posttest 2.)

Note that these four situations are likely to have very different plasticity demands, and it is not known which is the true purpose of adaptation. Second, even if that were known, the precise needs of each situation are not clear. For instance, injuries are so varied that being able to isolate one small region of space could be beneficial. Finally and most important, even if we knew the precise plasticity needs, that does not guarantee a perfect match between those needs and the actual abilities. Suppose that the purpose of flexibility is to accommodate different environments, and it could be further determined that five distinct types of mappings are needed. With more than one or two mappings needed, a mechanism may have evolved to decompose all of them into individual pairs rather than handle each type separately. As an unintended consequence, many other mappings could be accommodated even if they will never be encountered in the real world.

The empirical finding that mappings are not represented in this way suggests that prior research on elementary learning processes have only limited applicability. This is significant because the theoretical concept of a simple associative connection continues to influence learning research. This is evident not only in current explanatory mechanisms of animal learning (e.g., Rescorla, 1992) but in current models of human cognition (e.g., Gluck & Bower, 1988; Rumelhart, Hinton, & McClelland, 1986). In addition, principles of elementary associative learning continue to be applied to phenomena of perceptual learning (e.g., Hall, 1991). Although small regions along the visual continuum are perceptually discriminable, they were found not to be independent. Consequently, rather than representing a mapping as a collection of independent visual-proprioceptive pairs ($v_1 p_1, v_2 p_2, \dots, v_n, p_n$), the interrelation is better captured by the following general formula: $P = f(V)$, where $f(V)$ refers to a specific class of functions, yet to be identified, with a few modifiable parameters. The form of the function reflects its internal structure, and learning is reflected by a change in parameter values of that function (see also Hay, 1974). I refer to this as *dimension learning*.

Dimension Learning Contrasted With Single-Stimulus Learning

When general cognitive abilities were used, adjacent regions were easily separated. In this situation, it is more plausible to suggest that the fundamental building block is a single connection between a (discriminable) stimulus and a response. I refer to this type of process as *single-stimulus learning*. Single-stimulus and dimension learning can be viewed as two points along a single continuum that differ in the size of the region that may be independently manipulated. For one dimension of space, the size of that region could in principle have ranged anywhere from a few fractions of a degree to approximately a 180° semicircle. Although by this analysis single-stimulus and dimension learning would differ only quantitatively, the two ends of the continuum nonetheless have qualitatively different consequences for learning and therefore should be considered distinct classes of processing. At least three differences follow from consideration

of independence of spatial positions: form of generalization, ease of learning, and explicit representation.

Form of generalization. If one location does not influence other locations, the expected form of generalization is the familiar decrement whereby a newly trained motor response rapidly dissipates for other untrained stimuli along the continuum. In dimension learning, all locations are dependent or linked to one another. Rather than the classic generalization gradient, the form of generalization will depend on the form of the built-in structure-function(s). If the function is linear, for instance, the generalized responses will have changes in slope and intercept only. The present study found that two-pair discrimination training led to the standard generalization gradient with the cognitive paradigm but not for perceptual learning.

Ease of learning. Second, if small regions are independent, then all mappings should in theory be both learnable and equivalent in difficulty. Mappings would be coded as a list of stimulus-response pairs, and each entry would be ignorant about the contents of another entry. Training even very different motor responses for nearby visual stimuli should not interfere with one another. In practice, one might expect performance limitations resulting from a long list or a large number of distinct responses to implement. Even with performance limitations, it should be easy to learn an isolation mapping in which only one small region of space is singled out for a response different from its neighbors. For dimension learning, ease of learning a mapping will depend on how closely that mapping conforms to the built-in function. Inconsistent mappings will be learned only with difficulty or not at all. Which mappings are readily learnable will provide clues as to the nature of the function. In the present study, the isolation mapping was easily mastered in the cognitive paradigm but poorly accommodated for perceptual learning.

Explicit representation. Finally, independence implies that each pair does not provide information about any other. Consequently, each trained pair must be represented explicitly. We expect to be able to distinguish between stimuli that were trained and those that were not. Tests might include comparing errors, reaction time, or variability between trained and untrained stimuli. In contrast, when pairs are not independent, the trained pairs themselves need not be remembered or individually represented. They are used to generate parameter values that in turn represent the relation between entire dimensions. One should not be able to discern from any measure of performance which stimuli were trained and which were not (see also Carroll, 1963; Koh & Meyer, 1989). Would the perceptual and cognitive data from the present study show differences on the criterion of explicit representation as well as the other criteria?

Figure 8 shows the between-subject variances as a function of position for both experiments in which two positions were trained (Experiments 1 and 4). The patterns clearly differ: The variability in the perceptual data remains low for stimuli far removed from the training stimuli, whereas the variability in the cognitive data grows with distance from the trained stimuli. The patterns reflect specifically that for perceptual learning—subject agreement is just as high for untrained stimuli as for trained stimuli—whereas for cognitive

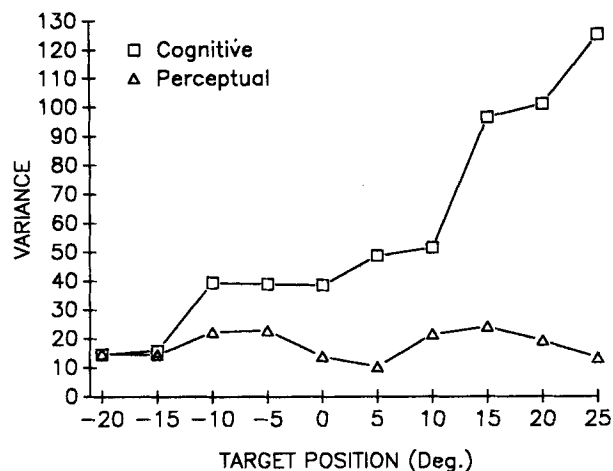


Figure 8. Between-subject variance as a function of target position for two-pair perceptual learning (Experiment 1) and two-pair cognitive learning (Experiment 4).

learning, agreement is high for explicitly trained locations only. Although these data are between-subject rather than within-subject variances, they do provide indirect evidence that trained and untrained stimuli are distinguishable in one process but not in the other.

Perceptual Learning

Why was the intuitively simple isolation mapping “only one region has changed” not learned? What is the built-in structure that precludes this mapping?

The felt-position-sense hypothesis. Can the findings be accounted for by the familiar hypothesis of a change in felt position sense?³ Harris (1965; see also 1980) suggested that the end product of adaptation was an alteration in the position sense of the limb seen through the prism. Erroneously pointing to a target following exposure to prisms was not due to seeing the target in a different place but to feeling the arm in a different place. Harris (1965) remarked that “according to the proprioceptive-change hypothesis, the subject comes to feel that his arm is where he saw it through prisms—even though this makes that arm’s position sense erroneous” (p. 421).

The proprioceptive-change hypothesis, however, is neutral as to whether the change is confined to specific trained regions or can influence the entire dimension. It cannot explain or predict what will happen for positions not explicitly trained. It does not dictate that the change must be a rigid shift for all spatial positions nor preclude different amounts of shifts for different positions. The hypothesis does not address which visual-proprioceptive mappings will be learnable and which will not. As Harris goes on to say, “It is not clear a priori whether a proprioceptive shift would make a subject misperceive arm positions other than those he saw while adapting” (p. 421). Rather, the hypothesis dictates that whatever the change, it affects the arm.

Work subsequent to Harris’s has suggested that another common locus of adaptation is within the eye such that the

felt direction of gaze becomes altered. Many studies have since been conducted on which training conditions lead to which endproducts, arguably the most prolific topic of inquiry on the phenomenon of prism adaptation (for reviews, see Howard, 1982; Welch, 1978, 1986). The present set of issues is orthogonal to that of where the change resides. Whereas the question of end product or locus of change asks where the change resides, the present issue of constraints asks what that change should be.

On the issue of *where*, Bockisch and Bedford (1990) found that the finger LED training paradigm produces a change that is on average approximately 28% visual, 68% (hand-arm) proprioception, and 4% unknown. To assess the components separately, several Harris-type tasks were used during testing in addition to pointing to targets. Changes in pointing straight ahead and in pointing to sounds were calculated for the estimate of proprioceptive changes, and changes in pointing to targets with the unexposed arm and selection of straight ahead visually were used for the estimate of visual changes. But as noted, breakdown into components cannot explain patterns of generalization or limitations on what mappings are learned, and explanatory principles of a different sort are necessary. In the terminology used by Marr (1982), constraints are needed for the level of the “computational theory,” whereas issues of the site of adaptation would be found in the level of “hardware implementation” and perhaps “representation and algorithm.”

Linearity. One possibility is that only linear functions, $P = aV + b$, can be used to realign the two spaces. Linear relations allow for both rigid shifts (intercept parameter b) such as in classic prism adaptation ($b =$ approximately 11°) and uniform magnifications and minifications (slope parameter a). In the present study, the nonlinear isolation mapping produced instead a linear change in behavior (Experiment 2). The rigid shift found would be the optimal solution if rules permit only a linear function. In addition, my prior work (Bedford, 1989) found that under a variety of training conditions with two pairs, interpolation was always linear.⁴

Topology. One of the isolation mappings (Experiment 2, “all but one region has changed”) did produce small nonlinear changes in pointing. That finding, along with the finding that extrapolative regions appear to be handled differently from interpolative regions (Experiment 1; see Bedford, 1989, in press), suggests that a linearity rule may not be a complete description of the constraints on realigning spatial dimensions. A second possibility is that the isolation mapping was particularly difficult because it was a nontopological transformation.

If the visual dimension of space is imagined as a rubber band, then different mappings can be visualized as different

³ Thanks to reviewer Robert Welch for suggesting that the issue of felt limb changes be explicitly discussed.

⁴ Some very recent work in related areas of adaptation have also found linear functions. Adaptation to rearranged auditory space produces a linear fit to a nonlinear transformation (when tested with distortion removed; Shinn-Cunningham, Durlach, & Held, 1993); phoria adaptation leads to a linear interpolation when trained with two points (Schor, Gleason, Maxwell, & Lunn, 1993).

physical transformations performed on the rubber band to realign it with proprioceptive space. For one-dimensional space, a transformation that preserves topological properties allows any type of nonuniform pulling and squishing of the rubber band, but a nontopological transformation also allows the rubber band to be actually cut or separate points to be glued together. From the viewpoint of transformation geometry (see Klein, 1893/1957; Modenov & Parkhomenko, 1957), these would be the most radical types of transformations possible because not even the original order of the points or the distinctness of points is preserved.⁵ In addition, Shepard (1989), while referring to neural representations and models, suggested 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).

For the "only one region has changed" mapping, the visual locations (-15° , -10° , -5° , 0° , 5° , 10° , 15°) were mapped into the proprioceptive locations (-15° , -10° , -5° , 5° , 10° , 15°). Note that the order of points is not preserved: Going from left to right, the position 0 used to be encountered after -5 and before 5; now, it is encountered after 5. The distinctness of points is also destroyed: Both 0 and 10 are mapped onto the single location 10, producing a many-to-one mapping. Therefore, the transformation is not only nonlinear but also not topological. It remains to be determined how mappings will be accommodated that are nonlinear but do preserve topological properties.

Hierarchy of rules. Transformation geometry may also explain the preference for linear mappings. In this view, there is a hierarchy of rules of which linearity is a preferred subset. Briefly, the complexity of any mapping can be determined by considering how many geometric properties of an initial form are altered by the transformation. Mappings that are categorized by this taxonomy as simplest would be accommodated the easiest, and so forth. Whereas transformations that do not preserve topological properties would be the most difficult, linear changes would be at the opposite extreme. For instance, rigid shifts can be imagined as sliding over the whole rubber band to a new location. This transformation is considered the simplest because nothing about the form itself is altered. Klein's hierarchy of geometries and transformations has been applied to other areas involving spatial relations, including the rat's conception of space (Cheng & Gallistel, 1984) and perception of facial growth (e.g., Mark, Todd, & Shaw, 1981) and may be applicable to this domain as well (see Bedford, in press).

Perceptual Learning Contrasted With Cognitive Learning

Presence of a discrepancy. What determines which learning system will be used? Perceptual learning may be generally described as occurring whenever, as a result of experience, the identical proximal stimulus leads to a different percept from before and continues to do so in the absence of new information. For instance, the same pattern of proprioceptive joint information may lead to a new conclusion about where the arm is, or the same efferent eye

signals interpreted as a different eye position from before. Why should this occur? Whereas the purpose of cognitive learning is to apprehend new information about the world, the purpose of perceptual learning is to correct internal malfunctions or to otherwise improve on the ability to perceive. The latter processes keep the sensory systems in good working order to allow the former processes to operate. But how will perceptual systems know that information contained in the proximal stimulus reflects an internal error rather than what is out in the world? One clue is the detection of a discrepancy of the sort discussed in the introduction. For instance, if the visual and motor systems detect that the same object is suddenly in two different places, a priori knowledge that this is an impossible state of the world will lead to the conclusion that it must be an internal error. If such a discrepancy is present, the motivation for an actual perceptual change is present. Otherwise, various cognitive abilities can be used to master a new task.

Awareness. A second factor concerns the role of awareness. Typically, awareness is considered unessential for perceptual learning. For instance, prism adaptation can take place without any knowledge that the visual image was transformed (cf. Welch, 1978, p. 19). On the other hand, awareness is usually thought of as defining a cognitive task. The distinction, however, may not be clear-cut. Awareness of the transformation can be present during perceptual learning. It tends to increase the rate at which subjects execute correct responses, although how it affects the underlying learning process is not well known (Welch, 1978). In addition, several cognitive learning and memory processes are now thought to occur without awareness or knowledge of the task (e.g., Schacter, 1987). Consequently, the role of awareness in distinguishing perceptual and cognitive learning remains unclear at present. (For discussion of the related issue of attention and prism adaptation, see Redding, Clark, & Wallace, 1985.)

Cognitive studies that use dimensions. Can cognitive processes operate on entire dimensions as well? Two studies that had subjects intentionally learn multiple stimulus-response connections (Carroll, 1963; Koh & Meyer, 1989, 1991) without any apparent perceptual discrepancy found evidence for dimension learning. Carroll trained subjects on 13 stimulus-response pairs that were generated by either an underlying linear function relating the two continua or a quadratic function, or were random connections. The stimulus was always the character *V*, which appeared in different positions from left to right within a horizontal band. The other continuum was also left-to-right location, indicated by a vertical line within another horizontal band underneath the first. During a testing trial, subjects were shown a stimulus in one location and were required to draw a vertical line in the appropriate location. Carroll found that subjects reproduced the correct locations more accurately for the linear and quadratic

⁵ Because a transformation that does not preserve topological properties preserves nothing of the original form, it typically is not considered geometry. Within the actual hierarchy of transformations and accompanying geometries, topological transformations are the most radical.

arrangements than for the random arrangement (and linear better than quadratic), suggesting that subjects were extracting an underlying function to perform the task. Koh and Meyer used visual-motor dimensions of length (the stimulus) and keypress duration (the response). Subjects were trained on eight length-duration pairs that conformed to different functions (power, linear, or logarithmic) and subsequently were tested on four new lengths in addition to the original stimuli. They discovered that both errors and variability were no higher for untrained stimuli than for trained stimuli. Like Carroll, they concluded that subjects had extracted the underlying relation and used that function to generalize behavior.

It is likely that for a cognitive learning task, many different approaches may be used, including breaking down a dimension into small independent units and treating a dimension in its entirety. The critical factors that elicit one or the other representation are not yet known. In the present study, two factors may have encouraged the majority of subjects to use a single-stimulus representation. Unlike the studies discussed earlier, only two distinct pointing responses were required, one located a few inches to one side of a target and one directly at the target. With few distinct responses, explicit storage of each trained pair is a sensible strategy. Second, instructions emphasized repeatedly to learn where to position the hand for each light in space, thereby calling attention to separate manipulability of individual regions.

In general, single-stimulus representation has the advantage of flexibility, enabling virtually any patterning between stimuli and responses to be learned. This flexibility is achieved at a cost of efficiency and speed because the greater the number of stimuli and responses, the greater the number of connections that must be explicitly trained. On the other hand, dimension learning is efficient; only a few pairs are needed to calculate the parameter values and infer the appropriate response everywhere. But flexibility is sacrificed: Patterns not consistent with the built-in structure cannot be accommodated. Note that when perceptual learning is needed, so far the evidence suggests that only dimension learning can be used. Perhaps it is essential for perceptual processing, particularly for a growing infant, that coherent perception and appropriate behavior be restored with only limited encounters with the seen and felt world (see Spelke, 1990).

Universality

Will the linkage of any two dimensions have the same built-in constraints? Carroll's findings are consistent with the present data: Linear functions are preferred. But Koh and Meyer's do not appear to be. They found that linear functions were not learned particularly well, whereas power functions were. The power functions observed, however, may be caused by mapping the physical dimension of duration onto the psychological dimension of duration (cf. Stevens, 1957). If true, the preferred relation between the psychological dimensions of duration and length may actually be linear. As Koh and Meyer point out, log transforms produce a linear function. Further investigation of dimension learning with

different dimensions and mappings is likely to prove fruitful for understanding learning mechanisms in general and perceptual learning in particular.

References

- 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. (in press). Perceptual learning. *Psychology of Learning and Motivation: Advances in Research and Theory*, 30.
- Berkeley, G. (1910). *A new theory of vision*. London: Dent. (Original work published 1709)
- Bockisch, C. J., & Bedford, F. L. (1990, April). *Locus of constraints on modifying spatial maps*. Paper presented at the annual meeting of the Rocky Mountain Psychological Association, Tucson, AZ.
- Carroll, J. D. (1963). *Functional learning: The learning of continuous functional mappings relating stimulus and response continua* (Rep. No. RB-63-26). Princeton, NJ: Educational Testing Service.
- 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: Erlbaum.
- Ebenholtz, S. M. (1974). The possible role of eye-muscle potentiation in several forms of prism adaptation. *Perception*, 3, 477-485.
- Epstein, W. (1975). Recalibration by pairing: A process of perceptual learning. *Perception*, 4, 59-72.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 14, 227-247.
- Hall, G. (1991). *Perceptual and associative learning* (Oxford Psychology Series, No. 18). Oxford, England: Clarendon Press.
- 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 reading and adaptability* (pp. 95-149). Hillsdale, NJ: Erlbaum.
- Hay, J. C. (1974). Motor-transformation learning. *Perception*, 3, 487-496.
- Hay, J. C., Langdon, B., & Pick, H. L., Jr. (1971). Spatial parameters of eye-hand adaptation to optical distortion. *Journal of Experimental Psychology*, 91, 11-17.
- Held, R. M. (1965). Plasticity in sensory-motor systems. *Scientific American*, 213, 84-94.
- Howard, I. P. (1982). *Human visual orientation*. Chichester, England: Wiley.
- 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. Olson & E. Smith (Eds.), *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society* (pp. 333-340). Hillsdale, NJ: Erlbaum.
- Koh, K., & Meyer, D. E. (1991). Function learning: Induction of continuous stimulus-response relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 811-836.
- Macintosh, N. J. (1974). *The psychology of animal learning*. London: Academic Press.
- Mark, L. S., Todd, J. T., & Shaw, R. E. (1981). Perception of growth: A geometric analysis of how different styles of change are distinguished. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 855-868.

- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Modenov, P. S., & Parkhomenko, A. S. (1965). *Geometric transformations, Vols. 1 and 2*. New York: Academic Press.
- Paap, K. R., & Ebenholtz, S. M. (1976). Perceptual consequences of potentiation in the extraocular muscles: An alternative explanation for adaptation to wedge prisms. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 457-468.
- Redding, G. M., Clark, S. E., & Wallace, B. (1985). Attention and prism adaptation. *Cognitive Psychology*, 17, 1-25.
- Rescorla, R. A. (1992). Hierarchical associative relations in Pavlovian conditioning and instrumental training. *Current Directions in Psychological Science*, 1, 66-70.
- 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.
- Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing* (Vol. 1, pp. 45-76). Cambridge, MA: MIT Press.
- Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 501-518.
- Schor, C., Gleason, G., Maxwell, J., & Lunn, R. (1993). Spatial aspects of vertical phoria adaptation. *Vision Research*, 33, 73-84.
- 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* (pp. 251-275). Cambridge, MA: MIT Press.
- Shepard, R. N. (1989). Internal representation of universal regularities: A challenge for connectionism. In L. Nadel, P. Culicover, L. A. Cooper, & R. M. Harnish (Eds.), *Neural connections, mental computations* (pp. 104-134). Cambridge, MA: MIT Press.
- Shepard, R. N. (1991, February). *Can natural selection yield universal principles of mind?* Paper presented at the meeting of the American Association for the Advancement of Science, Washington, DC.
- 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.
- Shinn-Cunningham, B. G., Durlach, N. I., & Held, R. M. (1993). *Adaptation to transformed auditory localization cues in a hybrid real-virtual environment* (Technical Report). Cambridge: Massachusetts Institute of Technology, Research Laboratory of Electronics, Virtual Environment and Teleoperator Research Consortium.
- Spelke, E. (1990). Origins of visual knowledge. In D. N. Osherson, S. M. Kosslyn, & J. M. Hollerbach (Eds.), *An invitation to cognitive science* (Vol. 2, pp. 99-127). Cambridge, MA: MIT Press.
- Stevens, S. S. (1957). On the psychophysical law. *Psychological Review*, 64, 153-181.
- Sutherland, N. S., & Macintosh, N. J. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.
- Taub, E. (1968). Prism compensation as a learning phenomenon: A phylogenetic perspective. In S. J. Freedman (Ed.), *The neuro-psychology of spatially oriented behavior* (pp. 77-106). Homewood, IL: Dorsey Press.
- Wallach, H. (1976). *On perception*. New York: Quadrangle.
- Wallach, H. (1984). Learning in space perception. In L. Spillmann & B. R. Wooten (Eds.), *Sensory experience, adaptation, and perception: Festschrift for Ivo Kohler* (pp. 353-364). Hillsdale, NJ: Erlbaum.
- Wallach, H., & Karsh, E. B. (1963a). The modification of stereoscopic depth-perception and the kinetic depth-effect. *American Journal of Psychology*, 76, 429-435.
- Wallach, H., & Karsh, E. B. (1963b). Why the modification of stereoscopic depth-perception is so rapid. *American Journal of Psychology*, 76, 413-420.
- 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. 24-1-24-45). New York: Wiley-Interscience.

Received October 9, 1991

Revision received July 21, 1992

Accepted July 29, 1992 ■