In Gallistel, R. (Ed.) (2002), *Stevens' Handbook of Experimental Psychology, Vol. 3, Learning, Motivation, and Emotion, 3rd Edition*, NY: Wiley.

CHAPTER 7

Perceptual Learning

PHILIP J. KELLMAN

INTRODUCTION AND BACKGROUND

When we think of learning, several prototypical ideas come to mind: the encoding of an item in memory, the connecting of one idea to another, the connecting of a response to a stimulus, or the learning of a motor sequence or procedure. Less commonly considered, both in ordinary intuition and in research, is *perceptual learning*. Perceptual learning refers to experience-induced changes in the way information is extracted. A large and growing set of research results indicates that such changes are not only possible but pervasive in human information processing. On a full spectrum of tasks, from processing the most basic sensory discriminations to apprehending the most complex spatial and temporal patterns and relations, experience improves the pickup of information, often by orders of magnitude. These improvements affect almost all skilled behavior, form important foundations of higher cognitive processes (such as language; see Chap. 11, this volume), interact with other kinds of learning in important ways, and furnish one of the most important components of high-level expertise.

Preparation of this chapter was supported in part by research grant R01 EY13518-01 from the National Eye Institute and by a grant from the MURI Program at the Office of Naval Research. The author thanks Randy Gallistel and John Hummel for helpful discussions.

Yet perceptual learning is not well understood. What are the mechanisms that enable information-extraction systems to change their operation? Is there one basic process or several in perceptual learning? What are the conditions that lead to perceptual learning? How does this kind of learning relate to plasticity at various levels of the nervous system?

These questions have cycled in and out of scientific concern for more than a century. William James, in his *Principles of Psychology* (1890/1950), noted several examples of extraordinary perceptual skills and emphasized the importance of perceptual learning for expertise, including achievements that we often think of as motor skills. A flurry of research in the mid-1960s put perceptual learning firmly on the scientific map. The field owes a great debt to the work of Eleanor Gibson and her collaborators around this time, culminating in a classic review (E. Gibson, 1969). For about a decade or so afterward, few papers were published in perceptual learning with human adults, for reasons that are obscure. Perhaps the focus in this period on perceptual development in infancy occupied many of the relevant investigators. Another factor may have been de facto boundaries between different research communities. Perceptual learning has often been omitted from or poorly integrated with research on both animal learning and human cognition.

Beginning in the mid-1980s and still accelerating, there has been a new wave of interest in perceptual learning. Much interest has been sparked by findings at the lowest sensory levels of sensory systems, that is, improvements in basic sensory acuities previously assumed to be relatively fixed. Interest in linking these changes to phenomena of neural plasticity has also helped to spotlight, and inform, studies of perceptual learning.

Definitions of Perceptual Learning

Over the years, the phrase "perceptual learning" has been used to refer to various ideas. Some are restrictive in implicating a particular process or mechanism, or in labeling particular kinds of experimental effects. For example, some have suggested that we consider as perceptual learning only those effects that can be shown to be specific to low levels of the sensory nervous system. If subjects' improvements in an orientation discrimination task proves specific to the trained eye, the trained retinal position, or the trained orientation, these characteristics would argue against explanations in terms of high-level cognitive strategies. Hence, one can be on safe ground in calling these "perceptual" changes.

As we will see, research has revealed a variety of effects in perceptual learning, and the theoretical situation is still in flux. This situation suggests that we be more eclectic and functional regarding definitions. Although it is reasonable to seek criteria to distinguish perceptual learning from other types of learning, it is premature to limit the domain in advance of a better understanding of the processes involved. The case of specificity in low-level perceptual learning is instructive. As we see later, it turns out that the specificity of learning varies substantially with relatively minor alterations in learning procedures, and apparently low-level effects are modulated by higher-level factors. It is possible that rather than engaging wholly different processes with minor paradigm changes, we are discovering characteristics of perceptual learning processes that are multilevel and flexible. In addition, taking a broad view allows us to consider significant improvements in information extraction that do not involve the most basic sensory elements. Perception involves the extraction of structure from the environment by means of the senses (e.g., J. Gibson, 1966, 1979). This structure may be relational and complex. As with perception, perceptual learning may involve not merely low-level sensory coding but also apprehension of relatively abstract structure, such as relationships in time and space.

For a broad and functional definition, it is hard to improve on that given by E. Gibson (1969, p. 3):

"Perceptual learning then refers to an increase in the ability to extract information from the environment, as a result of experience and practice with stimulation coming from it."

As my purpose is not to be vague but inclusive, I explore a number of more specific ideas about perceptual improvement. These are set out next, and they comprise a tool kit for interpreting experimental evidence throughout the chapter. Eventually, these particular notions about how experience changes information extraction will attain sufficient clarity to help adjudicate questions about different learning processes and the relation of perceptual learning to other forms of learning.

Perceptual Learning and Perceptual Development

One other definitional matter is worth mentioning. Sometimes the phrase "perceptual learning" is used to refer broadly to the many and substantial changes in perceptual capacities that occur in infancy. This view was an especially snug fit to classical empiricist ideas about the nature of perception. Learning to construct reality must be high on the agenda of a new perceiver, if his or her innate endowment includes only the ability to have sensations, and meaningful perception of objects and events must be constructed from combining sensations (Locke, 1857; Titchener, 1896), or sensations and actions (e.g., Piaget, 1954). In such a scheme, all meaningful perception and most of perceptual development must be perceptual learning. The reason is that from this perspective, sometimes termed *enrichment* (J. Gibson & Gibson, 1955), the initial connections between stimulus variables and perceptual representations of environmental properties necessarily arise by learning.

Improved understanding of early human perception does not sustain the overall view that initial perceptual competencies are established by learning (for a review, see Kellman & Arterberry, 1998). Many appear prior to learning, and many others arrive via maturation after birth. Some experience-induced attunements and improvements in perception also occur early in life. For some abilities, such as pictorial depth perception, maturational and enrichment learning explanations still compete to explain the original connections between stimulus variables and meaningful perceptual representations.

Accordingly, in this review, we do not define perceptual development as synonymous with perceptual learning. We are concerned with learning processes that appear to occur throughout the life span, including infancy. Most of what we know about these processes to date comes from experiments outside the infancy period. As we will consider, however, some experiments with young infants are also beginning to shed light on the nature of the learning processes, as well as on their role in perceptual development.

Discovery and Fluency

Progress in understanding perceptual learning will come from refining our conceptions of particular processes and mechanisms. Many different ideas have been proposed. These ideas about how perception improves with practice fall into two general categories. Some involve *discovery:* how perceivers uncover, select, or amplify the particular information, features, or relations required for some discrimination or classification. The second category of change might be called *fluency*. Fluency effects involve changes, not in the content of information extracted, but in the ease of extraction. At its extreme, practice in information pickup has been argued to lead to *automaticity:* processing that is fast and relatively insensitive to attentional load (Schneider & Shiffrin, 1977).

Although the distinction between discovery and fluency is conceptually clear, some phenomena may be difficult to classify. First, discovery and fluency effects often arise together in learning. Second, the dependent variables associated with discovery and fluency are not fixed. Sensitivity (a measure of detection or discrimination ability in signal detection theory) would seem the best indicator of discovery of new bases of response, whereas improvement in speed of processing tends to indicate fluency. However, this mapping can be misleading. Not every paradigm in which perceivers become faster with practice implies an improvement in fluency. An improvement in speed, for example, may reflect discovery of a better basis for response. Likewise, under time constraints, sensitivity may improve because the same information can be extracted more quickly.

COMPONENTS OF PERCEPTUAL LEARNING

In what follows, I explore perceptual learning in terms of (a) ideas about processes and mechanisms and (b) illustrations of phenomena at different levels of sensation, perception, and cognition. The strategy will be as follows.

First, I set out a few ideas about processes and mechanisms. Later, I examine them in greater detail in connection with specific experimental findings. These will be addressed in a progression from early sensory sensitivities, including work in several senses, to work in middle vision (perception of contours, objects, and surfaces) through higher level vision. The review is selective; my aim is to consider a variety of explanatory ideas in perceptual learning in connection with useful examples from the research literature.

Some Discovery Processes

The most remarkable fact about perceptual learning is that it can lead to new bases of response. In an extreme case, an observer may appear not to encode or register a feature or relation that after practice becomes the basis for reliable classification. It is often said about wine-tasting skill that a novice may be unable to distinguish two different wines on any basis. As expertise grows, it may become obvious that one of the two previously indistinguishable wines is a prized delicacy, whereas the other is a cheap, barely drinkable embarrassment. (Suffice it to say, learning to discern such differences can have expensive consequences.)

There are several useful ways to think about the nature of discovery effects. How is it possible for perceptual systems to discover previously unnoticed information?

Sensitivity Change versus Noise Reduction

Perhaps the most basic question in this regard is whether it is possible to show rigorously that learning truly improves sensitivity. If perceptual learning improves sensitivity in detection or discrimination, different kinds of explanations are required than if it merely changes response biases or leads to the attachment of a label or response to an already-encoded attribute. Several investigators have examined perceptual learning to ask this question using procedures of signal detection theory. They have sought to learn whether sensitivity changes occur and how these might be more precisely characterized in terms of biases, signal enhancement, or noise reduction (Dosher & Lu, 1999; Gold, Bennett, & Sekuler, 1999). I examine these approaches next. The evidence indicates that perceptual learning involves true increases in sensitivity. Such findings certainly make perceptual learning worth explaining, but they do not in and of themselves furnish the explanations.

Selection and Differentiation

An influential notion of perceptual learning is E. Gibson's (1969) notion of perceptual learning as *differentiation* learning. Differentiation involves the selection of relevant features or relationships in stimulation that are useful in making particular classifications or discriminations. Unlike the attaching of significance to information (so-called enrichment learning), the function of differentiation learning is to allow selection of relevant information from among the abundance of available information, most of which may be irrelevant to a particular task. Gibson put forth the specific hypothesis that what is learned in perceptual learning are *distinguishing features:* those properties that make the difference in a particular task in which the observer must classify a stimulus as being one kind of thing or another. This notion may be contrasted with the idea that experience with certain objects leads us to form general structural descriptions. In the latter case, learning might lead to more detailed representations for all object attributes. On the distinguishing features hypothesis, learning will specifically affect key contrasts. Some evidence supports the idea that we do preferentially extract task-relevant dimensions of difference (e.g., Pick, 1965).

The distinguishing features hypothesis can be applied to learning at different levels of complexity. In complex classification tasks, this may involve some kind of search among complex relationships to find invariant bases of classification. Basic sensory tasks may include overlap in the sets of elementary analyzers activated by two similar stimuli. With discrimination practice, one relies more heavily on the most relevant analyzers for making the discrimination, whereas those that are activated equally by both categories may be suppressed.

In recent years it has been suggested that some aspects of differentiation learning—selection of relevant inputs and suppression of irrelevant ones—can be modeled using neural-style network learning models (Dosher & Lu, 1999; Goldstone, 1998; Poggio, Fahle, & Edelman, 1992). The stimulus is encoded as an input vector, that is, as values along a number of stimulus dimensions or an array of analyzers varying, for instance, in terms of sensitivity to retinal position, orientation, and spatial frequency. An output layer contains nodes that correspond to the different response categories for a task. There may be one or more hidden layers between the input and output layers. Nodes at one layer are connected to all nodes at the next layer. These connections pass activation along according to the weights of their connection (connection strength) and the activation of the nodes themselves. At the beginning of learning, all nodes at one layer are connected to all nodes in the succeeding layer with random weights. Across learning trials, weights change, either by back-propagation of an error signal (supervised learning) or by unsupervised learning schemes (e.g., Hebbian learning, in which weights increase between units activated at the same time). This kind of model can apply to selection of relevant analyzers for basic sensory tasks. It can also encompass some relations among features, most obviously conjunctive ones, in networks with hidden units. Conventional networks may be limited, however, in dealing with abstract or symbolic relations (see the section titled "Perceptual Learning of Abstract Relations").

Attentional Weighting

A particular hypothesis about improved selectivity in perceptual learning is the notion that selective attention guides the pickup of relevant information (and possibly suppresses irrelevant information). Attention can be allocated to particular dimensions, such as color, or to particular features, such as red. This notion that learning involves a selection of relevant dimensions and the connection of particular values on dimensions to behavioral outcomes has a long tradition in learning research on animals and humans (Lawrence, 1949; Trabasso & Bower, 1968). The attentional weighting hypothesis is formally similar to the hypothesis of selection and suppression of analyzers in sensory discrimination learning. The processes may differ in the involvement of explicit attentional or strategic components. Ahissar (1999) suggested that these higher and lower level selective processes interact to produce perceptual learning.

Discovery of New Relationships or Features

Some examples of perceptual learning appear to involve the discovery of features or relationships that were not initially encoded at all. To make clear what is intended here, compare two learning situations. First, consider problems in which learners must figure out which of two patterns fits in the experimenter-defined category (or which pattern leads to a reward in some task). Suppose that the two choices always have stimuli that are circles or triangles that are red or black and large or small. Across trials, the learner can test hypotheses about what determines the correct choice. Trabasso and Bower (1968) considered problems of this sort in detail, both in human and animal learning, and found that performance could be accurately modeled by learning processes in which trials altered two parameters: the relevance of a given dimension (e.g., shape) and the reward value of particular values on relevant dimensions (e.g., red connected to correct choices). Although learning in such paradigms involves important issues, the most commonly studied problems raise only minimally the issues of discovery of potentially relevant dimensions and features. The ways in which stimuli differ are salient and obvious from the start.

For contrast, consider a trainee in art appraisal. When confronted with an authentic van Gogh painting and a clever forgery, the trainee may not be able to detect any difference at all, nor indicate what dimensions might be relevant. The expert, on the other hand, may find the differences in brushstrokes obvious. In this example, the expert's enhanced sensitivity arguably involves noticing details to which the novice is oblivious. Learning must somehow allow unnoticed information to become salient or at least efficacious in guiding classification. Discovery of new information also seems to apply to higher order patterns and relationships, including abstract ones. A chess grandmaster, for example, may notice at a glance that white's position is lost, due not to an imminent loss of a piece, but to a structural defect that will take many moves to prove fatal. The novice may be completely blind to the structural information enabling the grandmaster's diagnosis. Here, the relevant information is relational and abstract, involving relations of shape, color, and spatial position.

Whether it involves fine detail or complex relations, discovery of information to which the observer has initially zero sensitivity (in a signal detection sense) is perhaps the most mysterious aspect of perceptual learning. For basic features, one possibility (elaborated later) is that experience leads to isolating particular sensory analyzers most relevant to a task. Initially, responses may be heavily influenced by analyzers that are activated by stimuli in both of two categories to be discriminated. These overlapping responses are weeded out with learning, leaving performance to depend on the analyzers that discriminate best. For higher level relational information, a possibility is that new sorts of information are synthesized by conjoining features that are initially encoded, a notion referred to as *chunking* (e.g., Chase & Simon, 1973) or unitization (e.g., Goldstone, 1998). A different idea is that expert classification depends sometimes on discovery of new, higher order invariants that underlie some classification (E. Gibson, 1969). Such invariants may involve relations that go beyond mere conjunction of already encoded information; the relevant relationships may involve structural relations of many kinds, including highly abstract information. Abstracting grammatical relations from speech signals in language learning is a good, albeit possibly special, example. In language, the realm of possibly relevant relations may be more constrained by specialized learning mechanisms than in the general case of perceptual learning.

How can higher order relations involving rich structure be discovered in perceptual learning? Another parallel to language may provide a clue. Novel sentences are routinely produced and comprehended in natural language use, presumably because they are synthesized from sets of basic elements and relations. It is possible that perceptual learning proceeds from a set of basic encodings and a set of operators that can connect basic features and properties to form relationships of higher order. For example, learning what kinds of things are squares may involve encoding edge lengths and relations among edge lengths as registered by equal/different operators (Kellman, Burke, & Hummel, 1999). Our current understanding of such learning processes is modest. Understanding phenomena in which perceptual learning appears to depend on the discovery or synthesis of new relations is an important challenge for researchers.

Remapping of Perceptual Dimensions

Another idea about perceptual learning is that the mapping between different stimulus dimensions, or between stimulus dimensions and perceptual representations, may be shifted by experience. In this category are experiments using rearranged optical stimulation, such as the shifting of visual directions laterally via prism goggles (Bedford, 1989; Harris, 1965) or, more radically, the inversion of scenes using inverting prisms (e.g., Kohler, 1964).

The remapping of perceptual dimensions is an important but special category of perceptual learning. Most often, remapping involves relations between two channels through which the same environmental property is perceived. For example, it is important that the felt position and visible position of one's arm correspond, because the world contains not a haptic space and a visual space, but space. Accordingly, remapping or recalibration occurs for perceptual inputs that are intrinsically linked in this manner. The obvious function of remapping is to maintain proper coordination among the senses and between perceptual and motor activity. One clear application of remapping processes involves changes that occur during growth and development. For example, the radial localization of sounds depends on time, phase, and intensity differences given to the two ears. The specific mapping between interaural differences and radial direction depends on the size of the head, which changes as a child grows. Remapping processes sensitive to discrepancies across modalities may serve to maintain sensory and motor coordination (Knudsen & Knudsen, 1985). For adults, the need for remapping processes is less obvious; nonetheless, the capability for remapping when adults are subjected to altered stimulus inputs is dramatic. Perhaps such phenomena imply some ongoing need for recalibration, even in adults.

As the focus of this review is primarily on learning processes that lead to improvements in the pickup of information, I will not do justice to the literature or issues on remapping processes. The section titled "Spatial Intervals" elaborates one example. For a more comprehensive discussion of the issues in remapping, see Bedford (1995).

Fluency Processes

Automaticity

A classic example of improved fluency as a result of perceptual learning is the work of Schneider and Shiffrin (1977). In a series of studies, subjects judged whether certain letters in a target set appeared at the corners of rectangular arrays. Attentional load was manipulated by varying the number of items in the target set and the number of items on each card in a series of frames. Early in learning, performance was highly load-sensitive, but with extensive training subjects came to perform the task equally well within a range of target set sizes and array sizes. These results led Shiffrin and Schneider (1977; Schneider & Shiffrin, 1977) to claim that a transition occurred from *controlled* to *automatic* processing.

Item Storage

A wide variety of evidence indicates that experience with particular items facilitates subsequent performance on those items in classification tasks. This effect occurs even in cases in which subjects have extracted a clear classification rule and in cases in which the familiarity is based on aspects that are irrelevant (for a review, see Goldstone, 1998). The effects of instance learning may diminish over days or weeks in comparison to effects of learning some rule or invariant (Posner & Keele, 1967).

The fluency improvements from item storage appear to lie at the margin of what we would label perceptual learning. Such improvements have been described as a form of "imprinting" in which the stored trace may be functionally described as a new "receptor" or "detector" (Goldstone, 1997). A reasonable alternative is that these improvements in fluency derive from the associative connection of a particular stimulus representation to a particular categorization response (Hall, 1991). A useful criterion may be whether the perceptual representation itself changes as learning progresses or whether learning consists of the connection of a given representation to responses or other representations (see Chaps. 1 and 2, this volume).

Unitization

In contrast to processes of differentiation that occur from experience, unitization refers to the combining or connecting of encoded features to create chunks or units that make perceptual classification faster or more accurate. Evidence suggests that stimulus features that co-occur tend to become encoded as units. Such a process has often been invoked to account for fluent processing of letters (e.g., LaBerge, 1973) and words (e.g., Salasoo, Shiffrin, & Feustel, 1985). Such chunking processes may also come into play for spatially separated entities, including separated line segments (Shiffrin, 1996) and complex spatial configurations in chess (Chase & Simon, 1973).

Several ideas have been proposed to account for unitization. One is that items or parts that are simultaneously activated in short-term memory become integrated units in long-term memory (Shiffrin & Schneider, 1977). Computational models using neural-style units have also used synchrony of activation as the basis for chunking (Mozer, 1991). Some physiological evidence suggests that training leads to the development of specific neural responses that depend on configural relations (Logothetis, Pauls, & Poggio, 1995).

Interaction of Fluency and Discovery Processes. Fluency and discovery processes may interact in the development of expertise. Writing in *Psychological Review* in 1899, Bryan and Harter proposed that automatizing the processing of basic information was a foundation for discovering higher order relationships. These investigators studied learning in the task of telegraphic receiving. When the measure of words (in Morse code) received per minute was plotted against weeks of practice, a typical, negatively accelerated learning curve appeared, reaching asymptote after some weeks. With continued practice, however, many subjects produced a new learning curve, rising from the plateau of the first. For some subjects, a third learning curve ultimately emerged after even more practice. Each learning curve raised performance to substantially higher levels than before.

What could account for this remarkable topography of learning? When Bryan and Harter asked their subjects to describe their activity at different points in learning, responses suggested that the information being processed differed considerably at different stages. Those on the first learning curve reported that they were concentrating on the way letters of English mapped onto the dots and dashes of Morse code. Those on the second learning curve reported that dots and dashes making letters had become automatic for them; now they were focusing on word structure. Finally, learners at the highest level reported that common words had become automatic; they were now focusing on message structure. To test these introspective reports, Bryan and Harter presented learners in the second phase with sequences of letters that did not make words. Under these conditions, performance returned to the asymptotic level of the first learning curve. When the most advanced learners were presented with sequences of words that did not make messages, their performance returned to the asymptotic levels of the second learning curve. These results confirmed the subjects' self-reports.

Although the robustness of the phenomenon of three separable learning curves in telegraphic receiving has been questioned (Keller, 1958), Bryan and Harter's (1899) ideas about improvement, as well as the tests indicating use of higher order structure by advanced learners, remain important. Specifically, they argued that discovery of structure is a limited-capacity process. Automatizing the processing of basic structure at one level frees attentional capacity to discover higher level structure, which can in turn be automatized, allowing discovery of even higher level information, and so on. This continuing cycle—discovering and automatizing of higher and higher levels of structure—may account for the seemingly magical levels of human expertise that sometimes arise from years of sustained experience, as in chess, mathematics, music, and science. Bryan and Harter's study offers one of the most intriguing suggestions about how discovery and fluency processes interact and complement each other. Their 1897 article ends with a memorable claim: "Automaticity is not genius, but it is the hands and feet of genius."

Cortical Plasticity

Underlying changes in discovery and fluency in perceptual learning are changes in neural circuitry. Although linking particular changes to particular information-processing functions remains a difficult challenge (Buonomano & Merzenich, 1998; Edeline, 1999), much evidence suggests that modifications of neural circuitry accompany perceptual learning. As research progresses, a number of the discovery and fluency processes enumerated earlier may turn out to be connected directly to types of neural changes. Some of these possibilities are considered later.

Perceptual Learning and Other Concepts of Learning

What is the relation of perceptual learning to other concepts of learning? A compelling answer probably awaits a more precise understanding of process and mechanism. One relationship that is not fully clear is the connection of perceptual learning to the notion of *implicit learning* (defined as learning without awareness). Many tasks used to study implicit learning are perceptual learning tasks. Consistent with implicit learning, in complex perceptual learning tasks (e.g., sorting of newborn chicks by sex, chess playing), experts are often unable to explain what stimulus relationships they are using in classification. Also, in some kinds of amnesic patients, explicit learning processes appear to be disrupted, whereas performance on perceptual and implicit tasks remains intact. Yet it is not at all clear that perceptual learning and implicit learning are synonymous. Some perceptual learning may in fact be explicit; in some pattern-classification learning subjects can indeed point out what information they are using. Moreover, neuropsychological data may suggest more than one type of implicit learning.

Many other questions exist regarding the boundaries and relations between perceptual learning and other notions of learning. One issue is whether some aspects of perceptual learning can be understood in terms of more familiar associative learning concepts (Hall, 1991). Another is the involvement of perceptual learning in what are usually regarded as motor skills. When a baseball batter successfully hits a pitch, perceptual learning may account for the skill that allowed him to detect early in its flight that the pitch was a curveball. However, perceptual differentiation of the feel of swinging the bat in various ways may also have been involved in learning the muscle commands that produced the smooth swing.

Other questions involve not so much boundaries between types of learning, but components of information processing that are common to, or analogous among, different forms of learning. One is automaticity. Some aspects of information extraction become automatic with practice, but the same appears to be true of habitual motor sequences and reasoning patterns. The notion of automatization seems to crosscut several forms of learning. Another family resemblance involves the conditions for perceptual learning. Perceptual skill seems unlikely to be subsumed by the dichotomy of declarative and procedural knowledge, yet the conditions under which perceptual learning occurs appear to have much in common with procedural learning (see the section titled "Conditions Affecting Perceptual Learning").

PERCEPTUAL LEARNING IN BASIC VISUAL DISCRIMINATIONS

Physical devices designed to detect energy or to make simple pattern discriminations will have fixed limits of sensitivity. An extraordinary fact about human perceptual learning is that many of our most basic sensory thresholds are modifiable by experience. In examining psychophysical evidence for this claim, I focus on vision. In recent years studies have shown that this conclusion is true for discriminations involving virtually all basic dimensions of early visual encoding, including, for example, orientation (Dosher & Lu, 1999; Shiu & Pashler, 1992; Vogels & Orban, 1985), motion direction (Ball & Sekuler, 1982) and stereoacuity (Fendick & Westheimer, 1983). I sample the literature selectively, highlighting studies that raise interesting issues or point toward explanatory mechanisms. Much of the work attempting to connect learning effects to specific sites of cortical plasticity has involved senses other than vision; I explore some of this work at the end of this section.

Vernier Acuity

Vernier acuity—the ability to detect deviations from collinearity of two lines—is a basic measure of visual resolution. It is often labeled as a *hyperacuity;* the term refers to the fact that sensitivity to a particular spatial difference is smaller than the diameter of individual photoreceptors. In Vernier acuity tasks, thresholds for reliable detection of misalignment may be 10 arc sec of visual angle, about a third of the diameter of a photoreceptor.

This level of precision is perhaps one reason that researchers have often considered basic sensory acuities to be relatively fixed operating limits. Remarkably, as research has revealed in recent years, these operating limits can be strikingly improved by training. Westheimer and McKee (1978) carried out an early training study using a Vernier task in foveal vision. After about 2,000 trials, subjects' thresholds decreased about 40% on average. These results have been replicated and extended for both foveal vision (Saarinen & Levi, 1995) and parafoveal vision (Beard, Levin, & Reich, 1995). The effects appear to be resilient; tests carried out 4 months later showed no decline in the improvement.

Curiously, little improvement was found in a variant of the Vernier task, the three-point Vernier task (Bennett & Westheimer, 1991). In the three-point task, two vertically aligned dots are presented, and the subject must judge whether a third dot midway between them is displaced to the left or the right of the imaginary line connecting the upper and lower dots. Despite practice for more than 10,000 trials, no reliable threshold changes were found. The difference in outcome from the standard Vernier task could be due to the difference between stimuli. Additionally, a procedural difference could be relevant. Bennett and Westheimer gave 300 practice trials prior to measuring learning effects. Learning could have occurred rapidly during these practice trials. Other reports suggest that Vernier learning occurs rather rapidly, with most of the effects attained within 300 or so trials (e.g., Fahle, Poggio, & Edelman, 1992). The three-dot results are nonetheless discrepant with the slower course of learning found by Westheimer and McKee (1978). Adding to the confusion, Fahle and Edelman (1993) did find a long-term learning effect for the three-dot acuity task. It is not clear how to resolve these discrepancies.

Learning effects on Vernier tasks are highly specific to the stimulus orientation used in training (e.g., Fahle et al., 1992). This fact has motivated explanations emphasizing changes at relatively early levels of the visual system. For example, Fahle et al. proposed that task-specific modules are set up based on retinal or early cortical inputs. Specifically, they suggested that the set of analyzer responses (photoreceptor outputs in their model but orientation-sensitive cells in a more plausible realization) to an individual stimulus are stored as a vector. Each such template is connected to a response output ("left" or "right" in a Vernier task), obtained through supervised learning. After storage of a number of such examples, new stimuli can be classified by comparison to the templates. The templates are used as *radial basis functions* in that each template's response to the new pattern is determined by a Gaussian function of its Euclidean distance from the pattern in the multidimensional space that encodes the templates. The responses of the several templates are linked by weights to the response categories. In the extreme case in which all stimuli match stored templates, this model amounts to a look-up table. Models of this type show learning effects of a magnitude similar to that shown for human subjects (Fahle et al., 1992). On the other hand, this kind of model seems an unlikely account of other features of perceptual learning in hyperacuity. One is the fact that learning can occur without explicit feedback (i.e., without supervised learning). The other is that radial basis function models predict specificity in learning in terms of retinal location, eye, and orientation, yet learning effects transfer at least partially across these dimensions (Fahle, Edelman, & Poggio, 1995).

Findings about specificity have also been examined in attempts to localize learning effects anatomically. Because only certain layers of cortical area V1 and earlier parts of the visual pathway have significant numbers of monocularly driven cells (cells sensitive to inputs from only one eye), one strategy has been to train in one eye and test in another. For Vernier acuity, results have been inconsistent (Fahle, 1994, 1995), except for a fairly clear effect that learning transfers across eyes when the inputs for both are given to the same hemiretina (Beard et al., 1995). (The left hemiretina—right visual field—of each eye sends its information to the left hemisphere of the cortex.) This result is consistent with the idea that learning effects involve binocular cells in the trained hemiretina. Results involving specificity of location within a single visual field have been inconsistent (Fahle, 1994, 1995).

Orientation Discrimination

Orientation tuning is a basic feature of early visual analyzers which first appears at the cortical level. Cells in V1, the earliest cortical visual area, tend to have small receptive fields sensitive to particular retinal locations, with clear orientation selectivity. Except in the earliest layers of V1, most cells are driven binocularly. These facts make orientation sensitivity especially interesting in efforts to connect learning effects to particular cortical loci. For example, a learning effect that was specific to orientation and to the trained eye would suggest the involvement of monocular V1 cells.

Orientation sensitivity has been shown to improve with practice (Shiu & Pashler, 1992; Vogels & Orban, 1985). In Shiu and Pashler's study, reliable improvement was shown over nine blocks of 44 trials each, all conducted at specific retinal locations. When lines appeared in new locations (either in the opposite hemifield or in the other quadrant of the same hemifield), little or no transfer of learning was observed. Learning was also specific to the orientations used. Such effects are consistent with learning mechanisms that are specific to particular retinal locations, perhaps orientation-selective cells in early cortical areas.

A separate experiment by Shiu and Pashler (1992) indicated that when subjects judged brightness differences in the same set of stimuli, they did not gain improved discrimination abilities for orientation. The latter result suggests the importance of cognitive set, attention, or active task engagement in perceptual learning. Such factors, however, would seem to involve much higher levels of the nervous system. Thus, even perceptual learning effects involving basic visual dimensions may depend on an interplay of higher and lower levels of processing.

Orientation and Visual Search

Studies in which visual search depends on an orientation difference between a target and other items in an array also show substantial practice effects (Ahissar & Hochstein, 1999; Fiorentini & Bernardi, 1980; Karni & Sagi, 1991, 1993). Karni and Sagi (1991, 1993) tested visual search for a set of three parallel oblique lines that could be aligned vertically or horizontally in an array of horizontal lines. The stimulus onset asynchrony (time between the display onset and a pattern mask) needed to achieve a given accuracy (e.g., 80% correct) decreased rapidly from the beginning of training. (After 3 sessions of about 1,000 trials each, it had decreased 50% for some subjects.) Improvements continued more slowly for many sessions afterward.

Some aspects of their data led Karni and Sagi (1991, 1993) to argue that learning consists of two components. One is a fast component that is noticeable within sessions. This learning fully transfers across eyes. The other component arises more slowly, appears to require some period of consolidation or sleep after learning (as discussed later), and is specific to the trained eye. Karni and Sagi reported that both kinds of learning are specific to the stimulus orientations used and the specific quadrant of the visual field. (Targets always appeared in the same quadrant during training.) The idea that different learning processes have both differing time courses and specificity characteristics is appealing for distinguishing underlying mechanisms. Unfortunately, it is not clear how consistently these effects occur. Schoups and Orban (1996), for example, used the same task as Karni and Sagi and found that learning of both types transferred fully across eyes.

Motion Perception

Discrimination of motion direction improves substantially with practice. Discriminating two directions differing by 3 deg is initially quite difficult but becomes highly accurate with extended practice. When training involves only a difficult discrimination, the effects are found to be largely specific to the training direction (e.g., Ball & Sekuler, 1982). Such specificity suggests alteration in the sensitivity of specific neural channels selective for that direction.

Results differ, however, if training utilizes an easier discrimination. Liu (1999) had observers perform a forced-choice discrimination between two directions differing by 8 deg. After training, performance improved at directions that differed by 90 deg from the training directions. Liu and Weinshall (2000) reported another interesting result. For direction discriminations involving stimuli 4 deg apart, there was little transfer when performance was tested at new orientations (90 deg different). A different measure of transfer, however, produced some evidence that learning does generalize. During the second discrimination task, the learning *rate* was almost twice as fast as that in the original task.

These results mandate some caution in inferring the neural locus of learning effects from specificity of learning effects. If the same learning processes are at work in the difficult and easier problems, there may be some differences in the way they are engaged by slightly different tasks. If so, varying levels of specificity may reflect more about the task than about the mechanism. Alternatively, the results may indicate different kinds of learning processes—one involving improved selectivity of particular neural channels and another utilizing higher level processes. These issues are considered further in the section titled "Task Difficulty" below.

Specificity in Perceptual Learning

At this point, it is fair to say that attempts to use anatomical or stimulus specificity to infer the locus of learning in the nervous system have not yielded any clear generalizations. This conclusion comes both from inconsistent results on nearly identical tasks and from differences across tasks. The situation may reflect the fact that multiple types of learning effects (involving multiple loci) are strongly affected by small task differences. An alternative conception is that perceptual learning ordinarily involves a coordinated interaction of higher level processes, such as attention, and lower level ones, such as tuning of receptors sensitive to particular stimulus properties (e.g., Ahissar & Hochstein, 1999).

Sensitivity versus Noise Reduction in Perceptual Learning

Applications of concepts of signal detection theory have led to recent progress in understanding perceptual learning. A fundamental question about changes in detection and discrimination performance is whether these effects entail true improvements in sensitivity. Sensitivity might increase, for example, if learning somehow amplifies the relevant internal signals used in a task. Another possible account of improved performance is that internal noise—departures from ideal processing within the observer—is somehow reduced. Dosher and Lu (1999) described a framework for studying these questions, illustrated in Figure 7.1.

The framework begins with the assumption that processing of a sensory discrimination may be thought of as assessing a signal's match to one or another internal templates. Matching accuracy may be affected by external noise (in the stimulus) or by internal noise. Internal noise may be of two types: additive (constant) or multiplicative with the stimulus (in which the stimulus equals the signal plus external noise). By testing performance with different amounts of external noise, different notions of improvement can be tested. In Figure 7.1a, the effect of practice is to enhance or amplify the stimulus; perhaps a better description is that calculation efficiency improves. Associated with this effect is a characteristic set of curves shown in Figure 7.1b. These curves show the signal contrast required to achieve a certain performance threshold. In

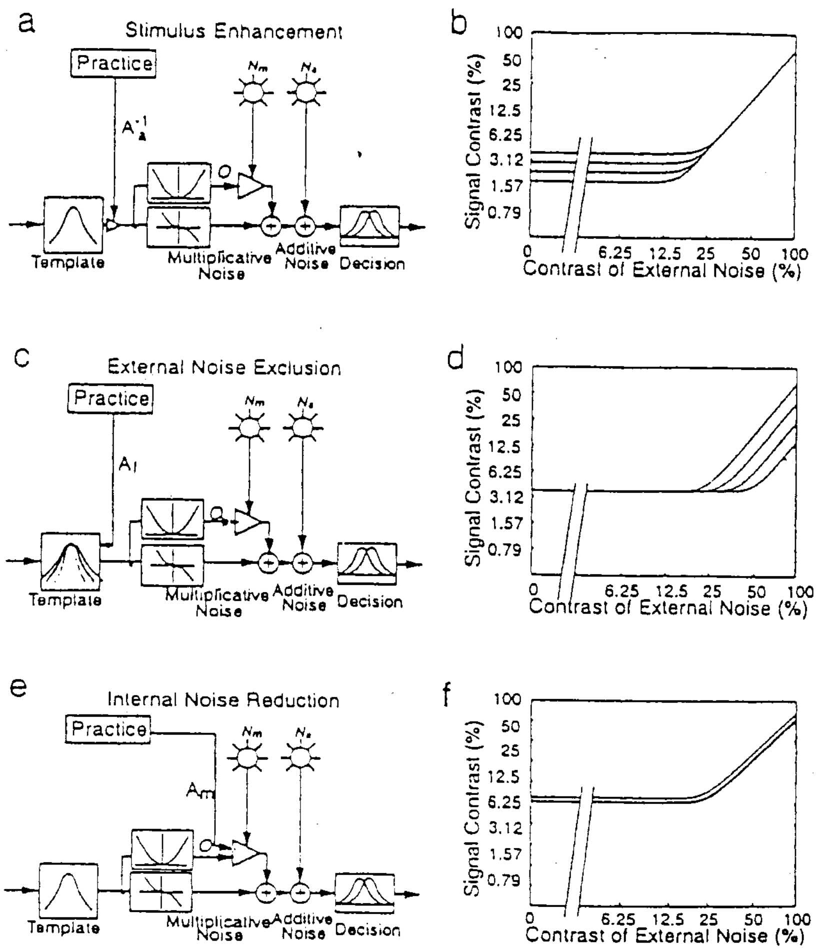

Figure 7.1 Models of learning effects and their data signatures.

NOTE: a) Practice turns up the gain on the stimulus, corresponding to stimulus enhancement. (N_m and N_a indicate multiplicative and additive noise, respectively; A_a indicates multipliers on internal additive noise, leading to stimulus enhancement.) b) Stimulus enhancement is associated with improvements in performance in the lower noise limb of the contrast threshold functions. c) Practice affects the amount of external noise processed through the perceptual template by narrowing the filter tuning, corresponding to external noise exclusion. (A_f indicates multipliers on the output of the perceptual filter applied to external noise, corresponding to external noise reduction.) d) External noise exclusion improves performance only in the high noise limb of the contrast threshold functions. e) Practice reduces the gain on multiplicative internal noise, or internal multiplicative noise reduction. (A_m indicates multipliers on internal multiplicative noise.) f) Internal (multiplicative) noise reduction improves performance somewhat over both limbs of the contrast threshold functions.

SOURCE: Reprinted from *Vision Research, 39,* B. A. Dosher and Z.-L. Lu, "Mechanisms of perceptual learning," pp. 3197–3221. Copyright © 1999, with permision from Elsevier Science.

the flat part of the curve, changes in external noise have little effect, as performance is limited by internal noise. Beyond a certain point, increases in external noise increase the contrast required to attain threshold performance. Practice has the effect of lowering contrast thresholds (shifting the performance curve downward) in the lower limb of the curve, while not affecting the higher limb.

Figure 7.1c illustrates the case in which learning reduces only external noise. (This is shown in the figure by the increasingly narrow tuning of the template.) The effect of practice on the data is schematized in Figure 7.1d. There is no change in the required contrast to attain threshold in the flat part of the curves. However, the range within which performance is limited by external noise (the rising portion of the curve) moves rightward with practice. Finally, the possibility of reduction of the multiplicative component (gain) of internal noise is shown in Figure 7.1e. Practice has the effect of improving performance in both parts of the curves, as shown in Figure 7.1f.

Dosher and Lu (1999) tested these predictions in a series of experiments in which observers judged on each trial whether a Gabor patch embedded in noise tilted to the left or right of vertical. A staircase procedure was used to find a contrast threshold in each condition. To distinguish different mixtures of the possible kinds of effects, experiments were carried out in which data were collected at more than one threshold level. This manipulation provided more detailed information on shifts in performance curves from practice.

Fitting the model predictions to their subjects' data, Dosher and Lu found evidence for both stimulus enhancement and external noise exclusion. There was no evidence of an effect of decreased gain of internal multiplicative noise. They discussed these effects in terms of changes with practice in the weights given to particular analyzers. Among a set of analyzers that are initially engaged by a task, some turn out to be less relevant than others. Learning is conceived of as a reduction of the contributions of less relevant analyzers to decisions.

One limitation in Dosher and Lu's analysis is the difficulty of distinguishing stimulus enhancement from reduction of additive internal noise, as these make similar predictions. Using similar methods, Gold, Bennett, and Sekuler (1999) found signal enhancement in perceptual learning tasks. They tested learning in a face identification and texture identification task. Subjects made a 10-alternative, forced-choice decision about which stimulus was presented on each trial in varying amounts of external noise. To assess the effect of internal noise, they used a double-pass response consistency measure, in which subjects judged a specific set of stimuli twice. Given that the signal plus external noise combinations in the stimulus set were held constant, any inconsistencies in responding must be due to internal variability. From the observer's consistency, it is possible to obtain an estimate of total internal noise (additive plus multiplicative). Gold et al.'s results suggested no change in internal noise.

Although the tasks in these studies have varied, all have indicated effects of signal enhancement and perhaps improved external noise exclusion. Changes in internal noise have not been found. It is not clear how general these findings are across different learning tasks.

Cortical Plasticity

The rapidity of perceptual learning in some studies is consistent with some known phenomena of neural changes in the brain. The activity of single neurons in sensory areas of the cerebral cortex is often characterized by their *receptive fields*—the description of the range of values on some relevant stimulus dimensions that influence the firing rate of the neuron. The receptive field of a visual neuron, for example, could describe the locations on the retina that, if stimulated, influence that cell's responding.

Visual receptive fields of cortical neurons can be changed rapidly by creating an artificial scotoma (blank area) in the cell's original receptive field while stimulation (e.g., moving gratings or dynamic noise dots) are presented in the surround. Pettet and Gilbert (1992) observed large increases in receptive field sizes within as little as 10 min of exposure to the artificial scotoma. It is not completely clear that increased receptive field size is the accurate description of the changes; an increased responsivity in the entire receptive field, including previously subthreshold areas, might explain the data (Das & Gilbert, 1995). Either kind of change might represent a cortical basis of perceptual learning. Contrary to behavioral studies of perceptual learning, however, such receptive field changes appear to be short-lived. It is possible that the difference in duration of the receptive field and perceptual learning effects are due to differences in training protocols.

A variety of studies have found evidence consistent with the specific idea of Hebbian learning mechanisms, in which co-occurring activations of units result in the strengthening of their connections. For example, the orientation sensitivity of V1 cells in cats can be experimentally altered by pairing presentation of selected orientations at the retina with applications to single cells of electrical current (Frégnac, Schulz, Thorpe, & Bienenstock, 1988) or neurotransmitter substances, such as GABA or glutamate (McLean & Palmer, 1998). Hebbian learning appears to be one mechanism of change in neural circuitry that could contribute to perceptual learning phenomena.

Plasticity in the Somatosensory System

Changes in cortical neurons, and indeed in whole cortical areas, are also characteristic of learning in other sensory modalities. For example, Wang, Merzenich, Sameshima, and Jenkins (1995) trained owl monkeys on a tactile task in which two bars were attached across three fingers at either their bases (proximal end) or their tips (distal end). Stimulation of all three fingers from one bar and the other generally alternated, and the monkey was required to respond whenever two consecutive stimuli were applied through the same bar. Normally, receptive fields in somatosensory area 3b are specific to individual digits; multidigit receptive fields are extremely rare. As shown in Figure 7.2, after prolonged training many cells exhibited multidigit receptive fields. The investigators noted that the development of these receptive fields is consistent with Hebbian learning.

Plasticity in the Auditory System

Whereas the organization of both early visual and somatosensory cortical areas involve topological maps of the space on the receptor surfaces (retina or skin), early auditory areas are *tonotopic*, organized in terms of frequency responses. The receptive field of a neuron in primary auditory cortex consists of the range of frequencies that influence its firing. Learning tasks have been shown to cause changes in the frequency responses of cells in auditory cortex of monkeys. Recanzone, Schreiner, and Merzenich (1993) showed that monkeys trained on a difficult frequency discrimination showed improvement over several weeks. Subsequent mapping of the receptive fields of cortical cells indicated that areas responding to frequencies relevant to the task were substantially enlarged. Weinberger and colleagues have found similar evidence of plasticity in guinea pigs and other species using a classical conditioning paradigm (Edeline & Weinberger, 1993; Weinberger, Javid, & Lepan, 1995). In most experiments a particular frequency was used as a conditioned stimulus (CS) and was paired with an electrical shock. Responses of cells in primary auditory cortex and also in the thalamus showed an enhancement of responses to the frequency used as the CS as well as a general alteration of many receptive fields such that they tended to become more centered on that frequency.

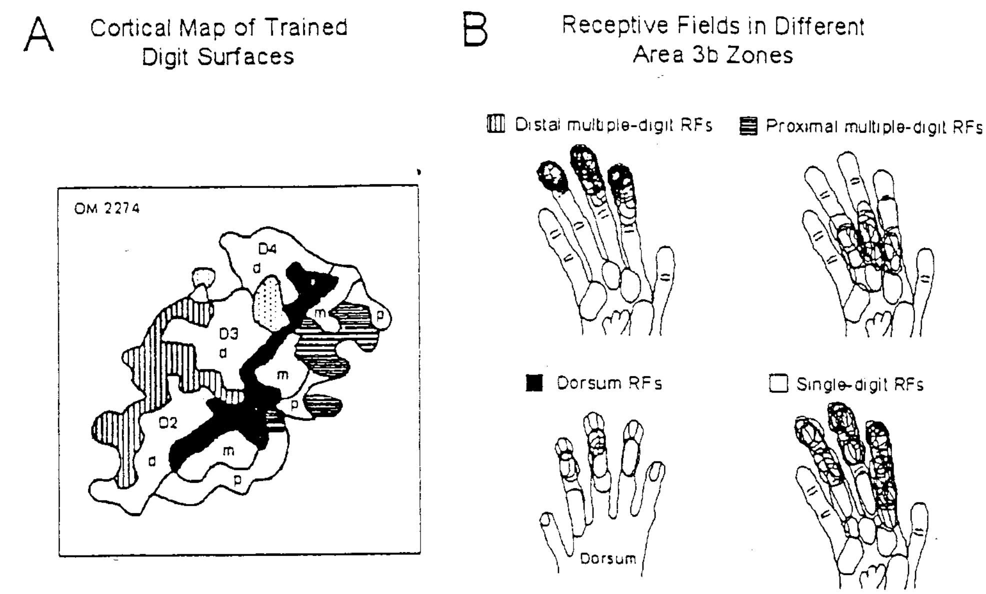

Figure 7.2 Training-dependent cortical map reorganization in primary somatosensory cortex. Map and receptive fields of the hand and representations of area 3b from the monkey that underwent behavioral training. Training involved extensive simultaneous stimulation across the proximal and distal portions of digits D2-D4. (A) The map shows that in contrast to normal maps, there was a significant portion of the map that exhibited multiple digit receptive fields, which were specific to either the proximal (horizontal striping) or distal (vertical striping) phalanges. (B) The receptive fields of the map shown in A, sorted according to the four observed classes; distal multiple-digit, proximal multiple-digit, dorsum, and single-digit receptive fields.
SOURCE: Buonomano and Merzenich (1998). Reprinted with permission.

Understanding the relations between particular kinds of neural plasticity and learning processes is a complicated task that will occupy researchers for a long time to come. One reason the task is so complicated is that the answers depend on other basic, unresolved issues, such as the precise nature of the learning processes themselves and how these processes and the information which they utilize are represented in the brain. An interesting discussion of specific issues confronting the effort to connect plasticity and learning phenomena may be found in Edeline (1999). One specific issue he raised is that we have as yet little understanding of how groups of neurons work together; yet clearly, circuitry encompassing more than the coding done by single cells is of fundamental importance.

PERCEPTUAL LEARNING IN MIDDLE VISION

Shape: Differentiation

A classic perceptual learning study (J. Gibson & Gibson, 1955) still serves as a good example of the idea of differentiation processes in perceptual learning. Figure 7.3 shows a coiled scribble pattern in the center, surrounded by a number of other patterns that differ along dimensions of compression, orientation, or left-right reversal. These patterns were combined with 12 others that were quite looked quite different from these scribbles and also varied greatly among themselves. In the experiments, cards containing individual patterns were shown. The standard stimulus (center

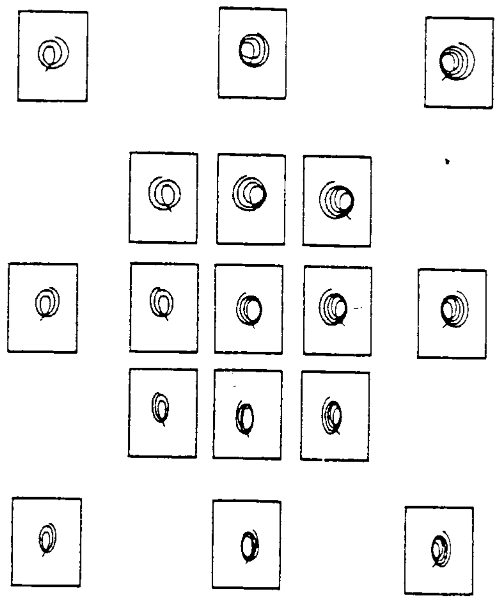

Figure 7.3 Scribble stimuli redrawn from Gibson and Gibson (1955).
NOTE: All patterns other than the one in the center are derived from the central pattern by compression, scaling, and/or left-right reversal. See text for additional details.

scribble) was shown for 5 s, and subjects were told that some items in the other cards to be presented would match the standard. They were instructed to make a same/different judgment about whether each card matched the previously shown standard. The dependent variable was the number of times through the pack of cards required to achieve perfect performance. Gibson and Gibson found that subjects tended to say "same" to similar items, but this tendency decreased across runs. Adults reached criterion (saying "same" only to the single matching stimulus) in fewer runs than did older (8.5- to 11-year-old) children, who learned faster than younger (6- to 8-year-old) children. Learning in this study proceeded without any explicit feedback. In general, errors involved overinclusion at the start, and exposure to the displays led to differentiation—the noticing of differences within the set.

Differentiation phenomena have a mysterious character. If, initially, the perceptual representations created by two or more stimuli are identical, how could they ever come to be discriminated? Perceptual representations of particular stimuli must somehow change with repeated exposure. Such changes may occur because of sampling or search processes. Of the many potentially encodable attributes of a pattern, only some are sampled and encoded in any one encounter (Trabasso & Bower, 1968). With repeated exposures, either some probabilistic element in the sampling process or a systematic search for dimensions of difference in a given task could lead to discovery of differences that were not initially noticed.

Whether initial stimulus encodings are followed by a search process that leads to differentiation may depend on the task in which the perceiver is engaged (E. Gibson, 1969). In the experiment with squiggles, the different patterns may initially have been encoded similarly as line drawings of roughly a certain size resembling coils of wire. In an experiment in which subjects were asked to judge members of this set to be the same or different, differentiation occurred. If the same patterns occurred incidentally on the sides of cars in a task where the observer was to classify the manufacturer of the car, these coils may have remained undifferentiated. The question of whether performance of an active classification is important in engaging perceptual learning mechanisms is discussed in the section titled "Active Classification, Attention, and Effort."

Shape: Unitization

The idea of unitization is that practice in a task allows features that were originally encoded separately to be combined into a larger unit. The term is synonymous with some uses of the term *chunking* (Chase & Simon, 1973). Goldstone (2000) reported a series of studies designed to examine unitization. He used

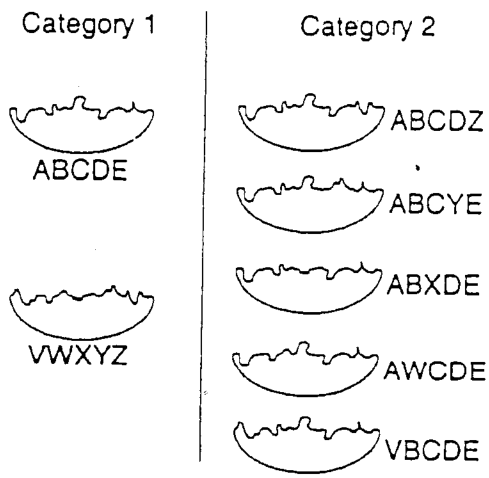

Figure 7.4 Stimuli used by Goldstone (2000).
NOTE: Each top contour is composed of five contour segments. See text for additional details.
SOURCE: Goldstone (2000). Copyright © 2000 by the American Psychological Association. Reprinted with permission.

stimuli like those shown in Figure 7.4. The top contours of these stimuli are composed of five separate shape fragments.

After extensive familiarization with the category assignments for the stimuli, subjects were given practice in sorting exemplars into categories 1 and 2 (shown in Figure 7.4). Reaction times were measured over four blocks of 80 trials. In one experiment, four conditions were compared. Categorization was based on (a) a conjunction of five parts in a particular order (as in Figure 7.4), (b) a conjunction of five parts with variable order, (c) a single component in a constant position, or (d) a single component in a randomly varying position. Responses showed that categorization based on a single component was faster than for those requiring multiple components. More important, the most learning was shown when categorization involved five parts in a consistent relationship. Goldstone interpreted the results to indicate that learners are developing a chunked representation that comes to function as a unit.

An important issue in Goldstone's studies of unitization, as in earlier work invoking unitization or chunking effects, is that there are at least two possible explanations for changes in performance. One is that a number of previously separate features come to be aggregated into a unit. The other is that processing of features gives way to categorization based on discovery of a higher-order relationship. The difference in these possibilities is that the higher-order relation is a new basis for response, not a collection of the lower level features. An example of a higher-order relation can be observed in display VWXYZ in Figure 7.4. If an imaginary curve connected the three highest peaks of the contour of VWXYZ, the curve would be nearly flat, perhaps slightly concave upward. In the other displays, such a curve would be convex upward. Thus, VWXYZ might be efficiently placed into category 1 based on this stimulus relationship. The relationship is not definable from the individual elements or from the mere fact that they occur together. This example is not meant to be a specific claim about the information that subjects used in Goldstone's studies; the point is that such higher-order relations are available. Goldstone explicitly indicated that his studies are unable to reveal the specific relations or units that the subjects actually use. His results appear consistent with either of these possibilities. One piece of evidence that true unitization is occurring is that learning effects also occurred in a condition in which the five contour fragments were separated and stacked in a vertical arrangement. The issue of stimulus redescription (discovery of higher-order relations) persists, however, as even a stack has invariant relations to be discovered. If discovery of higher order relationships is occurring, one might expect it to occur more readily for connected segments. Indeed, Goldstone's data indicate that learning was significantly better in the connected case.

Spatial Intervals

Perception of spatial extents is an important part of comprehending any environment. Spatial intervals can be signaled by a variety of information sources, and it appears that perceptual learning can function to maintain accurate perception from these sources. An example comes from research by Wallach, Moore, and Davidson (1963). They equipped observers with a telestereoscope, a viewing device that effectively changed the interocular distance (distance between the two eyes). In ordinary stereoscopic perception, the use of binocular disparity to specify depth involves a computation requiring the egocentric distance to at least one point in the scene. This distance to a point allows disparities to be converted into perceived depth intervals in the scene. Changing the interocular distance alters the magnitude of binocular disparities. When this occurs, the normal computation of depth intervals from distance and disparity is incorrect.

In the experiment of Wallach, Moore, and Davidson (1963), the observers viewed a rotating cube through the telestereoscope. Because all depth intervals were exaggerated, an edge of the cube that appeared to have a certain length when viewed horizontally appeared to grow in length as it rotated away to become more oriented in depth. With prolonged viewing, however, adaptation occurred, such that a new relation between disparity and depth obtained; this relation reestablished accurate perception of depth intervals. Wallach et al. argued that the basis of learning or adaptation in this situation is the assumption of physical invariance. The system adjusts the relation between depth and disparity (essentially learning a new interocular distance) in order to allow rotating object to have unchanging shape. A cogent elaboration of this type of argument, and further examples of this kind of perceptual recalibration, may be found in Bedford (1995).

As mentioned earlier, this kind of perceptual learning—remapping of the relationships across information sources—may differ in kind from most examples of perceptual learning that we have considered here. In adaptation research, including the large literature on adaptation to distorted optical input due to prism glasses and other devices, altered inputs lead to a new relation between perceptual dimensions, or between perceptual and motoric dimensions. Although involving multiple processes, a commonality of perceptual learning cases we have examined is that they involve improved information extraction with ordinary (unaltered) stimulation. Both kinds of learning are important.

Size Perception

Perceptual learning in size perception was tested by Goldstone (1994). In a pretest, subjects made forced choice, same/different judgments about the sizes of successively presented squares. Categorization training was then carried out in which four different sizes of squares had to be judged one at a time as "large" (the two largest squares) or "small" (the two smallest). After categorization, subjects were given a posttest that was identical to the pretest. They showed reliably improved sensitivity to size differences; enhancement was greatest around the category boundary. Some evidence suggests that these effects may be mediated by category labels and that they may not be truly perceptual effects, as evidenced by the lack of positive results when the same/different tests are carried out with simultaneously presented displays (Choplin, Huttenlocher, & Kellman, 2001). The roles of improved perceptual sensitivity, category labelling, and their possible interactions deserve further investigation, in this and other contexts.

Visual Search

As mentioned earlier, some researchers have argued that perceptual learning involves discovery of distinguishing features (E. Gibson, 1969), that is, those attributes or contrasts that govern some classification. An important question is how the learning of contrasts relates to the basic encodings that are unlearned. What kinds of features are naturally encoded prior to learning, and how can we tell? Treisman and colleagues (Treisman & Gelade, 1980; Treisman & Gormican, 1988) have attempted to characterize the basic inputs in vision as functionally separable *feature maps,* involving inputs such as orientation, size, color, closure of lines, and so on. Several criteria have been used to identify such features, which are said to be automatically encoded. One involves search times for an item having a certain feature in an array of items lacking that feature (distractors). If search times for a single item having the feature are insensitive to the number of distractors, that featural contrast is considered to be encoded in basic feature maps. Effortless segregation of textural regions based on the featural differences is another criterion. A third criterion relates to the converse idea: If information processing based on separately encoded features is relatively easy, processing of items that proves difficult may indicate that contrasts in single, automatically encoded features are not sufficient to do that task. One example involves conjunctions of basic features. Connecting information in separate feature maps to a unitary object is hypothesized to require attention. Accordingly, observers shown brief exposures of object arrays may experience *illusory conjunctions*—inaccurate conjunctions of features. Likewise, search for an item in an array whose difference from distractors is defined by a conjunction of features is slow, and it increases with the number of distractors (Treisman, 1991).

The idea that perceptual learning involves discovery or synthesis of new features suggests that this architecture may be modifiable. Some evidence indeed shows that the information-processing criteria used to identify basic features may be achievable for some classification tasks with practice. For example, Sireteanu and Rettenbach (1995) studied visual search tasks using feature contrasts that initially required serial search. These included searching for a plain circle among circles with gaps or with small intersecting line segments and searching for a pair of parallel edges among pairs of converging edges. Initially, performance indicated positive slopes relating reaction times to set size (1, 4, 8, or 16 items). A subject given extended practice achieved flat (approximately zero-slope) functions of set size for all search tasks. Some subjects who were tested for only a few hundred trials approached similar performance. A separate experiment showed that learning that achieved apparently parallel search on one task transferred fully to another search task and also transferred from the trained eye to the untrained eye. The tasks used in the transfer experiment involved searching for a circle target against distractors of circles with intersecting line segments and searching for a pair of parallel lines amid diverging line pairs. Transfer between these search tasks seems unlikely to depend on use of the same low-level analyzers. Instead, the results suggest that practice may lead to general search strategies that allow new feature contrasts to be processed as efficiently as those governed by basic feature maps.

The results of Schneider and Shiffrin (1977) may be related. Recall that in a task requiring search for target letters in a sequence of frames, each containing one or more letters, they found a transition from controlled to automatic processing, indicated by the fact that performance became insensitive to attentional load (number of target items times number of search items per frame). Shiffrin and Schneider (1977) interpreted their results in terms of connecting letter features in the target set to automatic attention responses.

The meaning of this account of improvement—attaching automatic attention responses to features—depends heavily on what can be a feature. One diagnostic criterion for identifying basic features, used by Treisman and others, has been asymmetry of search performance. Searching for a Q among Os allows "pop-out," because the system can distinguish, in theory, whether the feature map that registers straight line segments has some activation or none. Searching for an O among Qs, however, is slow and serial because both the feature map for closed loops and for segments have activation. To find the odd O requires attentive processes that examine particular locations. Taken at face value, the results of Sireteanu and Rettenbach (1995) appear difficult to fit with this overall scheme because practice led to parallel search for an O among what were essentially Qs. In this case, no feature is attached to an automatic attention response; rather, the location having the absence of the feature is what becomes easier to find. The results are similar to those involving conjunctive features (see the section titled "The Word Superiority Effect") in that they appear to require modification of our notions about basic input features or new ideas about what is learned with practice.

Implicit Learning in Visual Search

Earlier I noted that perceptual learning is sometimes considered to be closely related to the notion of implicit learning, that is, learning without conscious awareness. Implicit learning is often demonstrated by presenting certain stimulus regularities in the context of an irrelevant task and later testing to detect whether sensitivity to the regularities has been incidentally acquired (e.g., Reber, 1993). In recent years research has suggested that in certain populations of amnesics, explicit learning processes are impaired while implicit or perceptual learning processes are spared.

However, the precise relation between so-called implicit learning and perceptual learning remains unclear. It seems probable that multiple processes are encompassed by the conditions to which these terms have been applied. An example of the difficulty can be seen in research by Chun and Phelps (1999). They tested normal subjects and amnesic patients with hippocampal and surrounding medial temporal lobe system damage on a visual search task. The main experimental question involved the use of 12 stimulus arrays that repeated on half of the experimental trials. If subjects were able to encode the specific arrays, then over the 480 experimental trials, performance in locating the target should have become faster (because it always appeared in the same place in each of those 12 arrays). Results showed that both groups improved in overall visual search performance. However, only the normal subjects showed faster performance for repeated arrays after practice. Thus, what the authors call "contextual" learning in this paradigm appears to be impaired in amnesics, who can perform other implicit and perceptual learning tasks. Such results support neither a unitary notion of implicit learning nor an identification of perceptual learning with implicit (or explicit) learning.

PERCEPTUAL LEARNING IN HIGH-LEVEL PERCEPTION AND VISUAL COGNITION

Although some treatments of perceptual learning focus on changes in basic sensory function, even as a matter of definition, some of the most interesting and exciting phenomena and implications of perceptual learning involve relatively high-level information and tasks. These also pose some of the greatest challenges for understanding the processes and mechanisms involved.

Chess

One well studied example is expertise in chess. In 1997 the best human chess player, Gary Kasparov, defeated the best chess-playing computer, Deep Blue, in a 12-game match. There were some differences in the way the human and machine played. Deep Blue searched broadly and deeply through the space of possible moves and sequences at a rate of about 125 million moves per second. Skilled human players examine a smaller number of moves: about 4 per turn, with each followed about 4 plies deep (a ply is a pair of turns by white and black).

Given this difference in the scope of search, how could the human match the computer? The answer lies in human abilities to extract patterns from the board. These abilities were studied by DeGroot (1965, 1966), who tested both chess masters and less skilled players. Players at different levels differed not in terms of conceptual knowledge, search heuristics, or number of possible moves considered, but in their ability to encode and reconstruct a chess position after seeing it briefly. Grandmasters were able to reconstruct nearly perfectly a 25-piece position with a single 5-s exposure. This ability decreased substantially for players below the master level. It might be conjectured that chess masters and grandmasters happen to be individuals with exceptional visual memories, but that turns out not to be the case. When chess masters and ordinary individuals were tested for board positions that are not meaningful chess games, they showed equivalent performance in reconstructing the positions. It appears that chess masters and average individuals have about the same short-term memory capacities (Chase & Simon, 1973).

Chase and Simon (1973) hypothesized that experience with chess changes perception such that experts pick up chunks—sets of pieces in particular relations to each other. Following de Groot's work, they used a method in which subjects reconstructed viewed chess positions. Masters' overall performance (pieces reconstructed per view, number of views needed to reconstruct the whole position) was better than the overall performance of A-level chess players, which in turn exceeded that of beginners. Chase and Simon also measured the latencies with which subjects placed the pieces. Their data indicated that several related pieces—chunks—were placed in quick succession, followed by a pause and another set of related pieces, and so on. Masters had larger chunks than middle-level or novice players.

Chase and Simon concluded that most of the differences in chess skill related to changes in the way information is picked up that have occurred through practice. Their description involves both the concepts of discovery and fluency that we described earlier and is worth quoting:

> One key to understanding chess mastery, then, seems to lie in immediate perceptual processing, for it is here that the game is structured, and it is here in the static analysis that the good moves are generated for subsequent processing. What was once accomplished by slow, conscious deductive reasoning is now arrived at by fast, unconscious perceptual processing. It is no mistake of language for the chess master to say that he "sees" the right move; and it is for good reason that students of complex problem solving are interested in perceptual processes. (1973)

Similar comments may apply to almost any domain in which humans attain high levels of expertise. I use chess as an example partly because the value of perceptual learning can be quantified. Subsequent to his 1997 victory, Kasparov lost a close match with an improved Deep Blue that examined over 200 million moves per second. We can estimate that this human's ability to extract important pattern structure in chess is worth upwards of 125 million moves per second in raw search—an awesome equivalent computing power.

The Word Superiority Effect

Master-level chess skill is the province of very few; highlighting it as an example of complex perceptual learning may make such learning appear to be exotic and remote from ordinary cognition. The impression would be misleading. As an illustration, consider a phenomenon shown by almost every skilled reader of English: the word superiority effect.

Late in the 1960s, researchers discovered a remarkable fact. Exposure time required to identify which of two letters was presented on a trial was lower if the letter appeared in the context of a word than if the letter appeared alone (Reicher, 1969; Wheeler, 1970; see Baron, 1975, for a thorough review). In other words, subjects could more easily distinguish between WORK and WORD than they could judge whether a K or a D had been presented on a given trial. Detailed studies of the word superiority effect have revealed several other intriguing aspects. For one, the effect is not explained merely by rapid processing of familiar words as units. In fact, a substantial effect occurs for pronounceable nonsense (E. Gibson, J. Gibson, Pick, Osser, & Hammond, 1962). Baron and Thurston (1973), among others, found that pronounceable nonsense produced effects of the same magnitude as did actual words. These results indicate that general knowledge of some sort, perhaps the spelling or pronunciation patterns of English, facilitates letter recognition. One might predict that this kind of fluent processing of word-like strings would emerge from practice and skill at reading, and the prediction would be correct (Baron, 1978).

What kinds of mechanisms can explain the word superiority effect? Several detailed possibilities remain open (Baron, 1978; Noice & Hock, 1987). The effect may involve aspects of both discovery and fluency. Experience with English orthography seems to lead to detection of useful structures involving more than single letters. With practice these come to be rapidly extracted, so much so that the path from these higher structures to a decision about the presence of a particular letter ends up being faster than the time needed to detect the letter alone. The fact that the effect occurs for novel strings, not just for recurring words, suggests that the discovery of relations among letters generally characteristic of English spelling or pronunciation is involved. The fact that such units come to be processed rapidly and automatically provides a clue to the mechanisms of fluent reading. More specific understanding of what the relevant structures are and how they are learned remains to be discovered.

Unitization in Auditory and Speech Perception

Although this review has focused primarily on vision, processes of perceptual learning characterize information extraction in other sensory modalities as well. One conspicuous example, with some parallels to the word superiority effect in vision, is the learning of relationships in spoken words. Learning to segment the speech stream into words is a conceptually difficult problem, yet doing so is crucial for language learning (see Chap. 11, this volume). Saffran, Aslin, and Newport (1996) found evidence that learning can occur, even in 8-month-old infants, based on statistical relations among syllables. They used nonsense words comprised of three syllables and presented them in unbroken, monotone streams of 270 syllables per minute. Transitional probabilities between syllables X and Y (transitional probability = frequency of XY/frequency of X) were manipulated so that these were always high within words ($p = 1.0$ for syllables 1-2 and 2-3) and lower across "word" boundaries ($p = .33$). After familiarization for 2 min, infants were tested with sequences that either preserved ,words intact or changed the sequences of syllables. The showed a novelty response (longer attention) to the novel sequences. Follow-up studies (Aslin, Saffran, & Newport, 1998) showed that learning effects depend specifically on transition probabilities as opposed to more general frequency of exposure effects.

These findings provide dramatic evidence that forming of units based on statistical relations of sequential units is an early capacity of human learners, one that works in remarkably short periods of exposure. Is this capacity for early statistical learning specific to language learning? Saffran, Johnson, Aslin, and Newport (1999) explored this question by testing learning of statistical dependencies in tone sequences. Learning effects were similar, suggesting that these learning capacities may serve language acquisition but may also operate more generally.

Feature Conjunctions

A basic question in high-level perceptual learning is how new bases of response may be discovered. One natural source of such information is to make new combinations out of stimulus features that can already be encoded. Implicit in this idea is that perceptual learning may be something like a grammar—an open-ended class of new relations can be synthesized from a finite set of basic encodings and some means for combining these. Pursuing this general approach requires investigating both the vocabulary of basic encodings and the ways in which information can be combined.

Earlier I described some efforts to characterize features that are automatically encoded in vision (e.g., Treisman & Gelade, 1980). Feature contrasts that are basic (automatically encoded) may be the ones that allow efficient (e.g., load-insensitive or parallel processing) performance on certain kinds of tasks, such as visual search. Conversely, information that is not basic may require selective attention or serial processing. One prediction from this perspective is that search for conjunctions of features must utilize attention and must be sequential (across items) in nature. Numerous experimental tests have supported this conjecture. For example, Treisman and Gelade (1980) found that extensive experience did not reduce set size effects when subjects searched for a blue *O* among blue *T*s and red *O*s.

Yet the relatively fixed architecture that allows us to discover basic features (i.e., ease of processing features and difficulty with conjunctions) seems intuitively to be at variance with phenomena of perceptual learning and proficient performance. Various examples of expertise seem consistent with the idea that perceptual learning can lead to efficient processing of feature conjunctions. In chess, for example, knowing whether one piece is attacking another requires conjoining positions on the board, color, and shape. It is hard to imagine how grandmasters grasp whole board positions from a 2-s glance without being able to extract feature conjunctions efficiently, if not automatically.

The suspicion that feature conjunctions are learnable under some conditions is borne out by a small amount of experimental evidence. Shiffrin and Lightfoot (cited in Shiffrin, 1996) tested visual search for target patterns defined by conjunctions of spatially separated line segments and found that search slopes (average response time per element in the search arrays) decreased substantially with practice. Wang, Cavanagh, and Green (1994) found popout effects with the characters N or Z, shown in an array of backward Ns or Zs. The target was rapidly detected with little effect of the array size (number of distractors). The converse search task—searching for a backward N or Z in an array of Ns or Zs showed clear increases in response time with number of distractors. It appears that distractors can be rejected in parallel when these are highly familiar characters. The asymmetry in search performance suggests that the conjunctions of the several edges in the familiar characters have become encoded as unitary features (at least for purposes of rapid rejection in search).

Artificial Grammar Learning

Reber developed an important paradigm for testing the learning of structure: artificial grammars (Reber, 1967; for a review, see Reber, 1993). In his paradigm, letter strings were generated based on a grammar expressed as a transition network: Possible elements of the grammar (letters) were connected by directional paths, constraining the ways in which strings could be constructed. Although letter strings not permitted could look quite similar to those generated by a given grammar, evidence indicates that humans can, under certain conditions, learn the structural relations of the grammar allowing them to classify new strings correctly. Two issues that have been contested by researchers are whether the learning really consists of abstracting structure, as opposed to classifying new instances based on analogies to stored instances, and whether learning is implicit (i.e., without awareness). It appears that the learning can indeed consist of abstracting structure; it can also be based on analogies with stored instances, depending on the learning conditions (Reber & Allen, 1978). Evidence also clearly supports the idea that learning can be implicit. Structure can be detected from exposure to stimuli from a given generating grammar even when the structural relations are not directly relevant to an assigned task.

Implicit learning of artificial grammars is hardly ever discussed as an example of perceptual learning. Perhaps perceptual learning is too often interpreted as sensory learning (e.g., dealing with elementary sensory dimensions such as color) as opposed to the learning of structure in the input. Perhaps, because of their connection to linguistic material, artificial grammars may be thought of as too symbolic or abstract for perceptual learning. (For a discussion of these issues in the context of language acquisition, see Chap. 11, this volume). Recalling our earlier discussion of definitions, it is plausible to consider learning perceptual if it makes use of information in the stimulus. Structural relations in letter strings are fair game; if the learning arrives at an economical description (a grammar, for instance), that result might be better interpreted as indicating the nature of perceptual learning rather than indicating that the task is nonperceptual. Of course, some symbolic learning cannot be perceptual if the relevant information is not available in the stimulus. For example, in trigonometry one can learn from looking at graphs what a cosine function looks like, but one could not learn from looking that the function is ordinarily defined by a construction involving a triangle. In work with artificial grammars, there does not seem to be an extrastimulus component of this sort (i.e., the bases for learning are relations available in the stimuli). In a similar manner, although words have a primarily symbolic function, the fact that pronounceable nonsense strings show the "word" superiority effect implicates processes that pick up on stimulus relationships apart from the symbolic meaning in these kinds of representations. Integrating findings from implicit learning tasks used by cognitive psychologists with those in more commonly designated perceptual learning tasks may reveal commonalities and insights for modeling of processes that extract stimulus structure.

Object Recognition

Some evidence suggests that learning processes may specially engage object-specific representations. Furmanski and Engel (2000) measured exposure durations required for subjects to name low-contrast, gray-scale images of 60 common objects. Exposure durations needed to obtain 63% accuracy decreased 20% to 25% across five days of training. Transfer tests with a new object set showed partial transfer, indicating that effects involved both some generalizable learning and some specific component. A follow-up experiment indicated that learning for trained objects was fully maintained when half of them were mixed with an equal number of new object displays. Moreover, changes in size did not disrupt learning effects. These results differ from what would be predicted if subjects had learned primarily distinguishing features within the training set. The reason is that relevant contrasts for distinguishing objects should vary for different object sets. Results of this type suggest that perceptual learning may involve specific object descriptions as well as distinguishing features.

Perceptual Learning of Abstract Relations

An important characteristic of human perceptual learning is that it can involve abstract relations. Such an idea is consistent with theories of perception that emphasize abstract or higher order relationships (J. Gibson, 1979; Koffka, 1935) and the idea that perception produces abstract descriptions of reality (Marr, 1982), that is, descriptions of physical objects, shapes, spatial relations, and events, rather than, for instance, records of visual sensations. The idea that perceptual learning involves abstract relations is thus connected to the idea that perception itself involves abstract information and output.

What does it mean for perceptual learning to be abstract? It may help to provide a working definition, or at least a clear example, of what is abstract. Abstract information is information that necessarily involves *relations* among certain inputs, rather than collections of the inputs themselves. These ideas were central in Gestalt psychology (e.g., Koffka, 1935; Wertheimer, 1923). Their example of a melody serves as well today. Suppose you hear a melody and learn to recognize it. What is it that you have learned? At a concrete sensory level, the melody consists of a sequence of certain particular frequencies of sound. For most human listeners, however, learning the melody will not involve retaining the particular frequencies (or more properly, the sensed pitches corresponding to those frequencies). If you hear the same melody a day later, this time transposed into a different key, you will recognize it as the same melody. Your encoding of the melody involves *relations* among the pitches, not the particular pitches themselves. We often hear this fact mentioned in a disparaging light, namely, that few humans have "perfect pitch." In fact, the example makes a marvelous point about ordinary perceptual learning. The extraction and encoding of relations in the stimulus is fundamental. (Of course, musicians with perfect pitch are still better off in that they undoubtedly encode relations as well as particular pitches.)

The Gestalt psychologists argued that this response to patterns, rather than to concrete sensory elements, is pervasive in perception. In attaining the most important and behaviorally relevant descriptions of our environment, encoding relations is more crucial than is encoding sensory particulars (J. Gibson, 1979; Koffka, 1935; von Hornbostel, 1927). Whether this holds true depends on the task and environment, of course.

That encoding of abstract relations is a basic characteristic of human perceptual learning is suggested by recent research with human infants. Marcus, Vijayan, Bandi Rao, and Vishton (1999) familiarized 7-month-old infants with syllable sequences in which the first and last elements matched, such as "li na li" or "ga ti ga." Afterwards, infants showed a novelty response (longer attention) to a new string such as "wo fe fe" but showed less attention to a new string that fit the abstract pattern of earlier items, such as "wo fe wo." Similar results have been obtained in somewhat older infants (Gomez & Gerken, 1999). These findings indicate an early capacity for learning of abstract relationships. It is possible that these results are special in that they involve speech stimuli (Saffran & Griepentrog, 2001).

Mechanisms of Abstract Perceptual Learning

For a number of learning phenomena involving basic sensory dimensions, we considered how perceptual templates may be refined through practice. In a two-choice orientation discrimination, for example, the "template" for each orientation may consist of responses from a set of analyzers (e.g., orientation-sensitive units spanning some range of orientations, spatial frequencies, and positions). In a network-style learning model, these input units may be connected, perhaps through a middle layer of "hidden" units, with identification responses. Learning consists of the strengthening of weights of the most relevant analyzers with particular responses. This type of model is consistent with a wide range of results in perceptual learning (e.g., Dosher & Lu, 1999; Fahle et al., 1992; Goldstone, 1997). Such models are concrete in the sense that the output responses are determined by weighted combinations of the elements of the input vectors, consisting of the responses of analyzers at some point in the sensory pathway (e.g., V1 cells in the orientation example).

These concrete models are inadequate to explain more abstract examples of perceptual learning. Consider a simple example. In a concept learning experiment you are given the task of learning to classify letter strings as to whether they are in category A. You are told that the strings "VXV" and "DLD" are in the category. Additionally, you are told that the strings "ABC" and "MRH" are not in category A. Now, suppose you are tested with the string "KSK." Is this string in category A? In the many times I have done this demonstration, I have yet to encounter anyone who did not answer "yes" with high confidence.

There are several important implications of this and similar examples. First, humans readily discover structure in these items. Second, the structure is abstract. The classification of a novel item "KSK" does not match any of the training items in terms of the specific letters at each position. Rather, the learner apprehends a higher order structure: the relation that the first and last elements of the strings in category A are identical.

Standard neural network models are concrete in a way that makes them incapable of this kind of learning. Such a model would have an input layer, possibly one or more hidden layers, and an output layer. The nodes in each layer would be connected to all those in the prior and succeeding layers. Learning would change the weights of these connections between layers, such that certain inputs would lead to certain outputs. For our example, the input layer might have 26 possible nodes to be activated in position one of a three-letter string; there would also be 26 possible input activations for the middle position and 26 for the last position. Training on an example such as "DLD" would strengthen the connections between the letter D in first position and the output response "category A." The L and final D in their respective positions would also be weighted toward this categorization response.

The abstraction problem is simply that after training with the examples given, the network would know nothing about the test string

"KSK." The reason is that not one of those input letters had previously appeared in those positions (or any positions, in our example).

In an analysis of a related problem—a device that takes strings such as "1010" as inputs and gives an identical output ("1010" in this case)—Marcus (2001) concluded that they are not solvable by conventional statistical learning methods. Nonetheless, only a few examples are sufficient to allow human learners to learn this input-output rule.

Discovery of abstract, higher order relations seems to be a natural and important feature of human perceptual learning. A key challenge is to characterize the processes and mechanisms of this kind of perceptual learning. These problems are very general. They apply to the learning of shapes and relations in vision, to melodies, phonemes, and words in audition, and to a great many other things.

There has not been much work on learning abstract relations. One possibility is that such learning combines statistical learning processes with early extraction of relational information. Kellman et al., (1999) demonstrated the plausibility of such an approach in modeling the learning of shape categories for quadrilateral figures (e.g., square, rhombus, parallelogram, trapezoid). They assumed that these categorizations are not built into visual processors but must be learnable at least via supervised learning (as, e.g., when a child sees someone point to certain objects in the world and hears a word such as "square"). Learning, moreover, must draw only on inputs that are *generic*, that is, that are known or can be assumed to be available from ordinary visual processing. Specifically, they assumed that the visual system (a) can locate vertices or points of very high edge curvature in a figure, (b) can encode intervertex distances on the retina (or real distances when adequate information for size constancy is present), and (c) that the learning system contains operators that can compare extents for equality (and produce a graded response as departures from equality—e.g., of two extents—increase). Simulations showed that the model could learn almost all of the planar shape names tested from one or two examples. The network generalized its classifications to novel exemplars that were rotated, scaled, or distorted versions of the training exemplars. It also handled certain distortions in ways similar to human classifiers (e.g., tolerance for applying a term such as "square" to shapes with minor deviations given by unequal sides or misplacement of one vertex). Both the domain and the model are simplified in a number of respects, but this general approach of combining early relational recoding of inputs with later stages of connectionist learning processes may hold a great deal of promise.

CONDITIONS AFFECTING PERCEPTUAL LEARNING

Research to date has yielded useful clues about the processes and mechanisms of perceptual learning, but we by no means have a thorough understanding. Even in the absence of a complete understanding of process, however, we can say a fair amount about the conditions that lead to perceptual learning.

Contrast

Perceptual learning is facilitated by comparison of positive and negative instances of some category, or by contrasting instances that fit into differing categories. In E. Gibson's (1969) view, contrast is the very essence of this kind of learning: What is learned are distinguishing features—those attributes that govern the classifications that are important to the task. Hence, perceptual learning is often described as differentiation learning. The notion that differentiation learning is facilitated by presentation of negative instances is at least as old as Pavlov, who wrote,

> The question can now be discussed as to how the specialization of the conditioned reflex, or, in other words, the discrimination of external agencies, arises. Formerly we were inclined to think that this effect could be obtained by two different methods: the first method consisted in repeating the definite conditioned stimulus a great many times always accompanied by reinforcement, and the second method consisted in contrasting the single definite conditioned stimulus, which was always accompanied by reinforcement, with different neighboring stimuli which were never reinforced. At present, however, we are inclined to regard this second method as more probably the only efficacious one, since it was observed that no absolute differentiation was ever obtained by use of the first method, even though the stimulus was repeated with reinforcement over a thousand times. On the other hand, it was found that contrast by even a single unreinforced application of an allied stimulus, or by a number of single unreinforced applications of different members of a series of allied stimuli at infrequent intervals of days or weeks, led to a rapid development of differentiation. (Pavlov (1927), p. 117, cited in E. Gibson, 1969, p. 117).

One way of thinking about the effects of contrast is to view perceptual learning as a filtering process. In a given task a wealth of information may be available. Learning consists of selecting those features or relationships that are crucial for some classification (E. Gibson, 1969). In this process, those stimulus attributes that do not govern the classification must be rejected or filtered out. The presentation of negative instances (or members of an alternate category) allows decorrelation of the irrelevant attributes with the classification being made.

Task Difficulty

Specificity of perceptual learning depends on task difficulty in training. Ahissar and Hochstein (1997) undertook a systematic approach to this issue. The researchers used a visual search task in which subjects had to decide whether a line having a unique orientation appeared in an array of uniformly oriented background lines. Two dimensions of difficulty were manipulated. One was positional uncertainty: The unique orientation could occur anywhere in the array on positive trials, or, in an easier version, it could occur in only one of two locations in the array (indicated at the start of the experiment). The other dimension was the difference in orientation between the target and background lines; differences of 16, 30, and 90 deg were used. Figure 7.5 shows the conditions and data.

The task was considered to be easy either if the orientation difference was 90 deg or if the difference was 30 deg but targets were limited to two positions. Similarly, the task was considered difficult for 30-deg differences with all positions possible or for 16-deg differences and two possible positions.

These categorizations showed utility in predicting transfer data. Transfer to new orientations and new retinal positions was substantial for subjects in the easy conditions, but transfer was minimal for subjects trained in difficult conditions.

Learning easy examples first may lead not only to better transfer but also to better learning of hard cases. In fact, a single clear or easy trial can lead to rapid improvement in classification performance on difficult problems, a result termed the *Eureka effect* by Ahissar and Hochstein (1997). These investigators suggested that the effect indicates an interaction between high- and low-level mechanisms, with the high-level mechanisms directing the search for distinguishing information, followed by the attunement and selection of relevant low-level analyzers.

The connection between difficulty of training problems and specificity of learning has

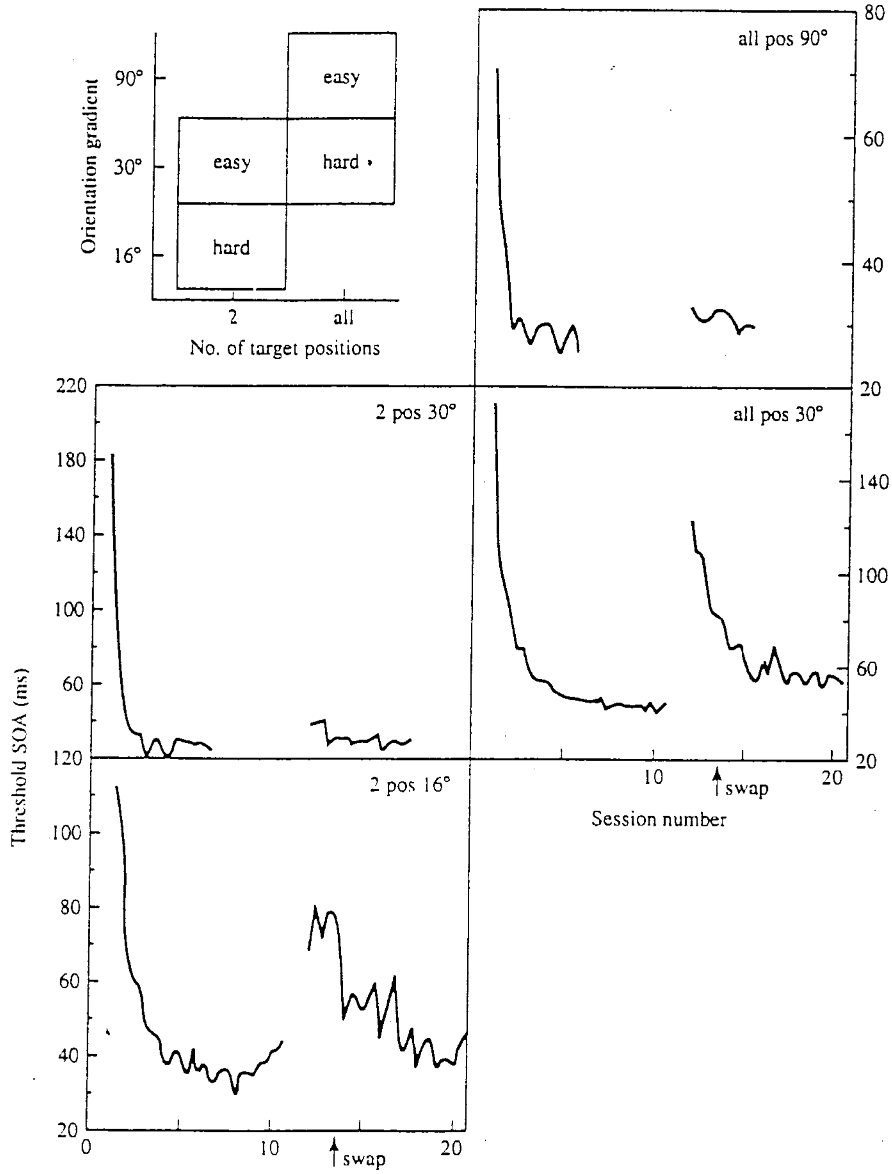

Figure 7.5 Learning data by condition difficulty in Ahissar and Hochstein (1999).
NOTE: SOA: Stimulus onset asynchrony. See text for additional details.
SOURCE: Ahissar and Hochstein (1999). Reprinted with permission.

been found in other learning tasks, for example, in Liu's (1999) studies of motion discrimination. Several investigators (Liu & Weinshall, 2000; Nakayama, Rubin, & Shapley, 1997), however, have suggested interpretations of difficulty effects—and rapid learning from clear examples—that do not involve communication between higher strategic processes and lower level analyzers. Instead, they propose that difficulty effects may reflect a single process. Recall that Liu and Weinshall (2000) found that learning a difficult motion discrimination did not transfer immediately to an orthogonal direction of motion, but it improved the learning rate for the second discrimination. Liu and Weinshall suggested that these results, and the results showing transfer across directions of easier discriminations, can be explained without invoking processes at multiple levels. They proposed that stimuli in a discrimination learning task initially activate populations of informative and uninformative analyzers. Because of computational capacity limits, on each trial the learning system samples only a subset of analyzers to assess their informativeness, an assumption proposed and supported in much earlier research (Trabasso & Bower, 1968). Over trials, learning effects of two types occur. Not only are individual analyzers are found to be informative or uninformative, but also *classes* of analyzers are assessed. Outputs of analyzers signaling particular spatial orientations, for example, may be irrelevant to a task in which motion direction must be discriminated.

These assumptions can be used to explain different kinds of learning and transfer effects (Weinshall & Liu, 2000). For difficult discriminations, only a few analyzers have high sensitivity for doing the task, and it takes longer to find these than in easy discriminations, in which many analyzers may have high sensitivity. Also, in the learning of one task, whole classes of analyzers may be discovered to be uninformative. Applying this logic to the transfer results for motion direction discrimination goes as follows. In learning the first problem, particular motion analyzers prove to be informative. These analyzers are not very helpful when a new direction of motion is tested. However, another effect of learning the first problem was that whole classes of analyzers (such as those for spatial orientation) have been learned to be uninformative. Therefore, particular members of such classes do not have to be sampled during the learning of the second problem; hence the faster learning rate for the second problem.

Consolidation and Sleep

Some evidence suggests that perceptual learning effects do not take hold until a consolidation or sleep period occurs. Data from Karni and Sagi (1993) suggest that learning effects from a particular session reach their peak after a consolidation period of about 8 hr. Stickgold, LaTanya, and Hobson (2000) used the same discrimination task as Karni and Sagi and found that maximal improvement occurred when subjects were tested 48 hr to 96 hr after a learning session. Their data also suggested the importance of sleep in consolidating perceptual learning effects. When subjects were deprived of sleep for 30 hr after the learning session, and then given two full nights of sleep recovery, they showed no improvement from initial levels.

The consolidation hypothesis has not been established beyond doubt. One alternative is that learning effects at the end of a learning session may be masked by fatigue effects. After a suitable interval, when fatigue effects have dissipated, learning effects are more visible (Shiu & Pashler, 1992).

Active Classification, Attention, and Effort

Studies in which the observer processes one stimulus dimension in an assigned task while stimulus variation on another dimension is simultaneously present find learning effects specific to the task-relevant dimension (Goldstone, 1994; Shiu & Pashler, 1992). Such results suggest the importance of attention in generating perceptual learning effects. Attentional effects were also suggested by Bennett and Westheimer (1991), who found no improvement in a grating acuity task tested in the fovea. One of their four subjects showed improvement for targets presented at 7.5 deg from the fovea, an effect they attributed to learning to make relatively large attentional shifts from the central visual field. Consistent with this idea, the reduction in threshold transferred fully from the horizontal training stimulus to a vertical stimulus. Unlike explanations invoking reweighting of low-level receptors, an attentional shift notion would predict transfer to a new orientation.

It is unfortunately difficult to untangle several conceptually different ideas here. One is that perceptual learning depends on attention. Another is that perceptual learning depends on the subject's active engagement in some kind of classification task. These possibilities may be hard to distinguish because any experimental manipulation that would involve assigning subjects an active task would also elicit their attention to the stimuli. The converse is easier to imagine: There are ways of arranging subjects' attention to stimuli without assigning a task. In some studies, mere exposure to certain stimulus variation during a task involving another stimulus dimension does not lead to learning (e.g., Shiu & Pashler, 1992). On the other hand, implicit learning of structure that is not specifically task-relevant is known to occur (e.g., Reber, 1993; Tolman, 1948). The inconsistency of results regarding learning from incidental exposure may be due to the possibility that carrying out some tasks involves active suppression of irrelevant information. In other words, learning incidentally while doing no task may be better than learning while doing some conflicting task. These issues of the roles of attention and assignment of active classification tasks in perceptual learning are ripe for further research.

Feedback

An intriguing characteristic of perceptual learning emphasized by E. Gibson (1969) is that in many cases it does not require feedback. Improved discrimination can come from mere exposure. In J. Gibson and Gibson's (1955) scribbles experiment (discussed earlier and shown in Figure 7.3), subjects judged whether a number of curved line patterns were the same or different from a sample pattern. If we describe the task as detecting differences, subjects initially made many errors that were misses (e.g., they labeled patterns that were physically different as "same"). Although no feedback was given, after several runs through the set of patterns, adult subjects achieved virtually perfect discrimination performance.

Models of statistical learning that work without feedback—unsupervised learning—may account for some aspects of exposure-based learning. In unsupervised schemes, the weights in a network change because of correlations in the inputs themselves (Hebbian learning), or they develop under certain constraints, such as the constraint that units in a hidden layer should be maximally uncorrelated with each other. A number of more sophisticated statistical techniques can be applied to unsupervised learning as well (for discussion of such methods applied to the problems of language acquisition, see Chap. 11, this volume).

Some studies indicate that perceptual learning of basic visual discriminations improves in similar fashion with and without feedback (Fahle & Edelman, 1993; Fahle et al., 1995; McKee & Westheimer, 1978). In Shiu and Pashler's (1992) study of learning in orientation discrimination, trial and block feedback (feedback only after each block of numerous trials) had similar effects on learning. A condition with no feedback showed smaller learning effects. Block feedback has also proved effective in other studies of perceptual learning (Herzog & Fahle, 1998; Kellman & Kaiser, 1994).

Herzog and Fahle (1998) pointed out that the effectiveness of block feedback, as well as a number of other commonly observed features of perceptual learning, are incompatible with conventional neural network architectures. They proposed a number of higher level (recurrent) mechanisms that guide learning via the allocation of attention and the use of feedback to modulate learning rates. One motivation for these suggestions consists of quantitative arguments indicating that learning occurs too efficiently to involve merely the gradual adjustment of weights in a network-style model. Higher level mechanisms that are sensitive to performance level and that perhaps incorporate knowledge of connectivity patterns in the nervous system may be implicated (Herzog & Fahle, 1993).

APPLICATIONS OF PERCEPTUAL LEARNING

There are numerous reasons to be interested in perceptual learning. As a somewhat neglected topic, its study adds new dimensionality to our ideas about learning. Many contemporary researchers view it, justifiably, as a window into processes of plasticity in the nervous system. As has been evident in this review, it is also a topic that connects various levels of information processing and neural activity in interesting and revealing ways.

Another reason for interest in perceptual learning is that it has great practical import. It is likely that changes in the way information is picked up—both in terms of discovery and fluency—form some of the most important foundations of human expertise. In an early section of her 1969 book, Eleanor Gibson included a section entitled "Perceptual Learning in Industry and Defense." Her examples included not only the grading of cheese and cloth, sexing of chicks, and wine tasting, but also some higher level skills such as landing an aircraft, interpreting maps and infrared photographs, and radiological diagnosis.

Although the grading of products, like William James' earlier comments about experts in Madeira and wheat, are not inconsequential examples, Gibson's other examples are perhaps of greater interest in indicating that the scope of perceptual learning applications may extend further into complex cognitive tasks than has generally been realized. Perceptual learning may underwrite abilities to discover complex, relational structures and become fluent in using them. The example of grandmasters in chess is instructive, as we saw. As argued by deGroot and by Chase and Simon, the most important distinguishing component of exceptional chess-playing ability involves learned skills for extracting patterns. It is not far-fetched to believe that such skills may be a major contributor also to the expertise of radiologists, pilots, financial analysts, mathematicians, and scientists.

Cognitive scientists and psychologists, however, have done much more to document the performance of experts than to apply perceptual learning concepts to develop expertise. In educational settings, one finds strong emphases on declarative facts and concepts and little attention to the development of expert apprehension of structure. One explanation may be the lack of any obvious method for bringing about expert information-extraction skills. For most advanced skills, from reading a financial spreadsheet to interpreting aircraft instruments, there are accepted methods of conveying facts and concepts, but the expert's intuitions about patterns and structure are believed to arise mysteriously from time and experience.

The view is too limited, however. Not only is the passage of time a suspect explanatory notion for perceptual skill, but also—strikingly—researchers have routinely been able to improve perceptual classifications in relatively brief laboratory experiments. As we have seen, these investigations have been carried out to address basic scientific questions, but they give hints as to methods by which advanced skills might be directly trained.

In recent years there have been attempts to apply perceptual learning methods to both basic and complex skills. Like experimental procedures, these ordinarily use training situations in which the subject receives many short classification trials. Successful efforts have been made to adapt auditory discrimination paradigms to address speech and language difficulties (Merzenich et al., 1996; Tallal, Merzenich, Miller, & Jenkins, 1998). For example, Tallal et al. reported that auditory discrimination training in language learning impaired children, using specially enhanced and extended speech signals, improved not only auditory discrimination performance but speech and language comprehension as well. Similar methods may be applied also to complex visual displays. Kellman and Kaiser (1994) designed perceptual learning methods to study pilots' classification of aircraft attitude (e.g., climbing, turning) from primary attitude displays (used by pilots to fly in instrument flight conditions). They found that an hour of training allowed novices to process displays as quickly and accurately as did civil aviators averaging 1,000 hours of flight time. Experienced pilots also showed substantial gains, paring 60% off their response times required for accurate classification. Studies of applications to the learning of structure in mathematics and science domains, such as the mapping between graphs and equations, apprehending molecular structure in chemistry, have also yielded successful results (Silva & Kellman, 1999; Wise, Kubose, Chang, Russell, & Kellman, 2000).

Efforts to improve directly the discovery of structure and its fluent processing are relatively new in educational and training contexts. Available evidence suggests that these have substantial promise, for both basic sensory discriminations and for processing of structure in complex and abstract cognitive domains. Much remains to be learned about the conditions that optimize learning, however. A number of lines of research have already suggested some of the conditions that affect the amount and durability of learning. To some extent, these can be investigated even in the absence of precise process models of perceptual learning. Understanding the variables that affect learning, of course, will have benefits beyond the practical. Clear accounts of when and how much learning occurs may be among the most important contributors to efforts to develop better models of process and mechanism.

SUMMARY

It has been more than 100 years since William James called attention to the phenomena of perceptual learning and over 30 years since the publication of Eleanor Gibson's synthesis of the field. What have we learned?

In several respects, progress is evident. More exacting tests of what changes in perceptual learning have been possible through the application of signal detection methods. Our tool kit of explanatory concepts has expanded and has also become more detailed in the form of computational modeling of notions such as analyzer weighting, differentiation, and unitization. At the level of biological mechanism, research is revealing types of plasticity that seem likely to relate to the implementation of perceptual learning processes in the brain.

However, most of these developments serve to sharpen our questions and to indicate how much remains to be learned. Accordingly, our answers to key questions, a few of which follow, must be necessarily provisional.

Is Perceptual Learning a Separable Form of Learning?

The evidence is persuasive that the general idea of improvements in the pickup of information deserves its own place among concepts of learning. Clearly, processes of information pickup do change with experience, and the representations that they produce change as well. These phenomena are not encompassed by other learning concepts, and their common involvement with information extraction allows them to form a natural grouping. At the margins are phenomena that may involve other forms of learning, such as associative or procedural learning. However, at the same margins, some phenomena thought to consist of the learning of procedures or connections between stimuli and responses no doubt involve changes in the way information is picked up and represented.

That said, clarification of the relationship between perceptual learning and other taxonomic categories of learning remains a high priority. The ways in which concepts such as implicit learning and automaticity crosscut several different forms of learning should be explored. Likewise, the relations among associative, procedural, and perceptual learning need to be further elaborated.

Is There One Process of Perceptual Learning, or Many?

The evidence seems clear that several processes are involved in perceptual learning. For example, the distinction between discovery and fluency processes may mark a difference in the kinds of improvement in perception. Discovering new bases of response may occur through the weighting of analyzers, the synthesis of new relation detectors, or the sampling of many information sources to locate the relevant ones. In all of these processes, changes occur in the content that is extracted. Improvements in fluency, on the other hand, can occur without changes in what is extracted; practice in particular information pickup tasks seems to increase speed and decrease attentional load and effort. Possible mechanisms include automaticity and unitization. At the borderline between discovery and fluency processes is the possibility that speed increases, not because of more rapid linking of existing representations but because of the discovery of higher order invariants that make classification more efficient. Distinguishing the operation of these processes is an important priority for research.

Does Perceptual Learning Involve a Single Level in the Nervous System, or Multiple Levels?

Both the existence of multiple processes in perceptual learning and our review of particular phenomena suggest that adaptive improvement in perceptual tasks involves multiple levels of neural activity. Despite the lack of consistency across tasks and procedures, in visual tasks with humans the existence of results indicating specificity of learning to a single eye, stimulus value, or retinal location strongly suggests involvement of cells at relatively early locations in the cortical visual streams. Physiological measurements of receptive fields in several senses and a variety of species directly implicate changes in other primary sensory cortices.

Meanwhile, a number of other findings indicate the involvement of higher level (including attentional and strategic) processes. In some paradigms, or with minor procedural changes from those showing specificity, learning does transfer: across eyes, across retinal positions, or across motion directions. Engagement of attention for a specific task affects learning, and initial presentation of easier examples facilitates it. Effects of trial-by-trial feedback do not indicate much about the locus of learning, but when block feedback produces effects on a par with trial feedback, the existence of higher level processes supervising and directing learning is intimated.

Finally, though not as often explored as yet by researchers, much of human perceptual learning involves abstract relationships. Little is known about the mechanisms of this sort of learning, but it cannot arise exclusively from modifications of receptive fields in primary sensory cortical areas. Attaining a computational and physiological account of high-level perceptual learning is among the many challenges remaining for researchers.

REFERENCES

Ahissar, M. (1999). Perceptual learning. *Current Directions in Psychological Science, 8*(4), 124–128.

Ahissar, M., & Hochstein, S. (1997). Task difficulty and learning specificity. *Nature, 387,* 401–406.

Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science, 9*(4), 321–324.

Ball, K., & Sekuler, R. (1982). A specific and enduring improvement in visual motion discrimination. *Science, 218,* 697–698

Baron, J. (1978). The word-superiority effect: Perceptual learning from reading. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes.* Hillsdale, N.J.: Erlbaum.

Baron, J., & Thurston, I. (1973). An analysis of the word-superiority effect. *Cognitive Psychology, 4*(2), 207–228.

Beard, B. L., Levi, D. M., & Reich, L. N. (1995). Perceptual learning in parafoveal vision. *Vision Research, 35*(12), 1679–1690.

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

Bedford, F. L. (1995). Constraints on perceptual learning: Objects and dimensions. *Cognition, 54,* 253–297.

Bennett, R. G., & Westheimer, G. (1991). The effect of training on visual alignment discrimination and grating resolution. *Perception & Psychophysics, 49*(6), 541–546.

Bryan, W. L., & Harter, N. (1899). Studies in the physiology and psychology of the telegraphic language. *Psychological Review, 6*(4), 345–375.

Buonomano, D. V., & Merzenich, M. M. (1998). Cortical plasticity: From synapses to maps. *Annual Review of Neuroscience, 21,* 149–186.

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology, 4*(1), 55–81.

Choplin, J., Huttenlocher, J., & Kellman, P. (2001, May). Perceptual discrimination and memory. Paper presented at the First Annual Meeting of the Vision Sciences Society, Sarasota, FL.

Chun, M. M., & Phelps, E. A. (1999). Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage [see comments]. *Nature Neuroscience, 2*(9), 844–847.

Das, A., & Gilbert, C. D. (1995) Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex. *Nature, 375*(6534), 780–784.

deGroot, A. D. (1965). *Thought and choice in chess.* The Hague, Netherlands: Mouton.

deGroot, A. D. (1966). Perception and memory vs. thought: Some old ideas and recent findings. In B. Kleinmuntz (Ed.), *Problem solving* (pp. 19–50). New York: Wiley.

Dosher, B. A., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. *Vision Research, 39,* 3197–3221.

Edeline, J. M. (1999). Learning-induced physiological plasticity in the thalamo-cortical sensory systems: A critical evaluation of receptive field plasticity, map changes and their potential mechanisms. *Progress in Neurobiology, 57*(2), 165–224.

Edeline, J. M., & Weinberger, N. M. (1993). Receptive field plasticity in the auditory cortex during frequency discrimination training: Selective retuning independent of task difficulty. *Behavioral Neuroscience, 107*(1), 82–103.

Fahle, M. (1994). Human pattern recognition: Parallel processing and perceptual learning. *Perception, 23*(4), 411–427.

Fahle, M. (1995). Perception of oppositely moving verniers and spatio-temporal interpolation. *Vision Research, 35*(7), 925–937.

Fahle, M., & Edelman, S. (1993). Long-term learning in vernier acuity: Effects of stimulus orientation, range and of feedback. *Vision Research, 33*(3), 397–412.

Fahle, M., Edelman, S., & Poggio, T. (1995). Fast perceptual learning in hyperacuity. *Vision Research, 35*(21), 3003–3013.

Fahle, M., Poggio, T., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science, 256*(5059), 1018–1021.

Fendick, M., & Westheimer, G. (1983). Effects of practice and the separation of test targets on foveal and peripheral stereoacuity. *Vision Research, 23*(2), 145–150.

Fiorentini, A., & Berardi, N. (1981). Learning in grating waveform discrimination: Specificity for orientation and spatial frequency. *Vision Research 21*, 1149–1158.

Frégnac, Y., Shulz, D., Thorpe, S., & Bienenstock, E. (1988). A cellular analogue of visual cortical plasticity. *Nature, 333*(6171), 367–370.

Furmanski, C. S., & Engel, S. A. (2000). Perceptual learning in object recognition: Object specificity and size invariance. *Vision Research, 40*(5), 473–484.

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

Gibson, E. J., Gibson, J. J., Pick, A. D., & Osser, H. A. (1962). A developmental study of the discrimination of letter-like forms. *Journal of Comparative & Physiological Psychology, 55*(6), 897–906.

Gibson, J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin.

Gibson, J. (1979). *The ecological approach to vision perception.* Boston: Houghton Mifflin.

Gibson, J. J., & Gibson, E. J. (1955). Perceptual learning: Differentiation or enrichment? *Psychological Review, 62*, 32–41.

Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Signal but not noise changes with perceptual learning. *Nature, 402*(11), 176–178.

Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General, 123*(2), 178–200.

Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology, 49*, 585–612.

Goldstone, R. L. (2000). Unitization during category learning. *Journal of Experimental Psychology: Human Perception & Performance, 26*(1), 86–112.

Gomez, R. L., & Gerken, L. (1999). Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition, 70*(2), 109–135.

Hall, G. (1991). *Perceptual and associative learning.* Oxford: Clarendon Press.

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

Herzog, M. H., & Fahle, M. (1998). Modeling perceptual learning: Difficulties and how they can be overcome. *Biological Cybernetics, 78*(2), 107–117.

Hornbostel, E. M. von. (1927). The unity of the senses (E. Koffka & W. Vinton, Trans.). *Psyche, 7*(28), 83–89.

James, W. (1950). *The principles of psychology.* New York: Dover Publications. (Original work published 1890)

Karni A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences, USA, 88*, 4966–4970.

Karni, A., & Sagi, D. (1993). The time course of learning a visual skill. *Nature, 365*(6443), 250–252.

Keller, F. S. (1958). The phantom plateau. *Journal of the Experimental Analysis of Behavior. 1*, 1–13.

Kellman, P. J., Arterberry, & M. E. (1998). *The cradle of knowledge: Development of perception in infancy.* Cambridge, MA: MIT Press.

Kellman, P. J., Burke, T., & Hummel, J. (1999). Modelling perceptual learning of abstract invariants. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the Twenty-First Annual Conference of the Cognitive Science Society* (pp. 264–269). Mahwah, NJ: Erlbaum.

Kellman, P. J., & Kaiser, M. K. (1994). Perceptual learning modules in flight training. *Proceedings of the 38th Annual Meeting of the Human Factors and Ergonomics Society* (pp. 1183–1187). Santa Monica, CA: Human Factors and Ergonomics Society.

Knudsen, E. I., & Knudsen, P. F. (1985). Vision guides the adjustment of auditory localization in young barn owls. *Science, 230*(4725), 545–548.

Kohler, Ivo. (1964). *The formation and transformation of the perceptual world.* New York: International Universities Press.

Koffka, K. (1935). *Principles of Gestalt psychology.* New York: Harcourt, Brace.

LaBerge, D. (1973). Attention and the measurement of perceptual learning. *Memory & Cognition, 1*(3), 268–276.

Lawrence, D. H. (1949). Acquired distinctiveness of cues, I: Transfer between discriminations on the basis of familiarity with the stimulus. *Journal of Experimental Psychology, 39,* 770–784.

Liu, Z. (1999). Perceptual learning in motion discriminatioin that generalizes across motion directions. *Proceedings of the National Academy of Sciences, USA, 96*(24), 14085–14087.

Liu, Z., & Weinshall, D. (2000). Mechanisms of generalization in perceptual learning. *Vision Research, 40*(1), 97–109.

Locke, John. (1857). *An essay on human understanding.* Philadelphia: Hayes & Zell.

Logothetis, N. K., Pauls, J., & Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biology, 5*(5), 552–563.

Marcus, G. F. (2001). *The algebraic mind.* Cambridge, MA: MIT Press.

Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science, 283*(5398), 77–80.

Marr, D. (1982). *Vision.* San Francisco: W. H. Freeman.

McKee, S. P., & Westheimer, G. (1978). Improvement in vernier acuity with practice. *Perception & Psychophysics, 24*(3), 258–262.

McLean, J., & Palmer, L. A. (1998). Plasticity of neuronal response properties in adult cat striate cortex. *Visual Neuroscience, 15*(1), 177–196.

Merzenich, M. M., Jenkins, W. M., Johnston, P., Schreiner, C., Miller, S., & Tallal, P. (1996). Temporal processing deficits of language-learning impaired children ameliorated by training. *Science, 271*(5245), 77–81.

Mozer, M. C. (1991). *The perception of multiple objects: A connectionist approach.* Cambridge, MA: MIT Press.

Nakayama, K., Rubin, N., & Shapley, R. (1997). Abrupt learning and retinal size specificity in illusory-contour perception. *Current Biology, 7*(7), 461–467.

Noice, H., & Hock, H. S. (1987). A word superiority effect with nonorthographic acronyms: Testing for unitizedvisual codes. *Perception and Psychophysics, 42*(5), 485–490.

Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex.* London: Oxford University Press.

Pettet, M. W., & Gilbert, C. D. (1992). Dynamic changes in receptive-field size in cat primary visual cortex. *Proceedings of the National Academy of Sciences, USA, 89*(17), 8366–8370.

Piaget, J. (1954). *The construction of reality in the child.* New York: Basic Books.

Pick, A. D. (1965). Improvement of visual and tactual form discrimination. *Journal of Experimental Psychology, 69*(4), 331–339.

Poggio, T., Fahle, M., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science, 256,* 1018–1021.

Posner, M. I., & Keele, S. W. (1967). Decay of visual information from a single letter. *Science, 158*(3797), 137–139.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning & Verbal Behavior, 6*(6), 855–863.

Reber, A. S. (1993). *Implicit learning and tacit knowledge.* New York: Oxford University Press.

Reber, A. S., & Allen, R. (1978). Analogic and abstraction strategies in synthetic grammar learning: A functionalist interpretation. *Cognition, 6*(3), 189–221.

Recanzone, G. H., Schreiner, C. E., & Merzenich, M. M. (1993). Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *Journal of Neuroscience, 13*(1), 87–103.

Reicher, G. M. (1969). Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology, 81*(2), 275–280.

Saarinen J., & Levi, D. M. (1995). Perceptual learning in vernier acauity: What is learned? *Vision Research, 35,* 519–527.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants [see comments]. *Science, 274*(5294), 1926–1928.

Saffran, J. R., & Griepentrog, G. J. (2001). Absolute pitch in infant auditory learning: Evidence for developmental reorganization. *Developmental Psychology, 37*(1), 74–85.

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition, 70*(1), 27–52.

Salasoo, A., Shiffrin, R. M., & Feustel, T. C. (1985). Building permanent memory codes: Codification and repetition effects in word identification. *Journal of Experimental Psychology: General, 114*(1), 50–77.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing, I: Detection, search, and attention. *Psychological Review, 84*(1), 1–66.

Schoups, A. A., & Orban, G. A. (1996). Interocular transfer in perceptual learning of a popout discrimination task. *Proceedings of the National Academy of Sciences, USA, 93,* 7358–7362.

Shiffrin, R. M. (1996). Laboratory experimentation on the genesis of expertise. In K. A. Ericsson (Ed.), *The road to excellence: The acquisition of expert performance in the arts and sciences, sports, and games.* Hillsdale, NJ: Erlbaum.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing, II: Perceptual learning, automatic attending and a general theory. *Psychological Review, 84,* 7, 127–190.

Shiu, L., & Pashler, H. (1992). Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Perception & Psychophysics, 52,* 582–588.

Silva, A., & Kellman, P. J. (1999). Perceptual learning in mathematics: The algebra-geometry connection. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society,* Mahwah, NJ: Erlbaum.

Sireteanu R., & Rettenbach R. (1995). Perceptual learning in visual search: Fast, enduring, but nonspecific. *Vision Research, 35,* 2037–2043.

Stickgold, R., LaTanya, J., & Hobson, J. A. (2000). Visual discrimination learning requires sleep after training. *Nature Neuroscience, 3*(12), 1237–1238.

Tallal, P., Merzenich, M., Miller, S., & Jenkins, W. (1998). Language learning impairment: Integrating research and remediation. *Scandinavian Journal of Psychology, 39*(3), 197–199.

Titchener, E. B. (1896). *An outline of psychology.* New York: Macmillan.

Trabasso, T., & Bower, G. (1968). *Attention in learning, theory, and research.* New York: Wiley.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review, 55,* 189–208.

Treisman, A. M. (1991). Search, similarity, and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception & Performance, 17*(3), 652–676.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12,* 97–136.

Treisman, A. M., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review, 95*(1), 15–48.

Vogels, R., & Orban, G. A. (1985). The effect of practice on the oblique effect in line orientation judgments. *Vision Research, 25*(11), 1679–1687.

Wallach, H., Moore, M. E., & Davidson, L. (1963). Modification of stereoscopic depth-perception. *American Journal of Psychology, 76*(2), 191–204.

Wang, Q., Cavanagh, P., & Green, M. (1994). Familiarity and pop-out in visual search. *Perception & Psychophysics, 56,* 495–500.

Wang, X., Merzenich, M. M., Sameshima, K., & Jenkins, W. M. (1995). Remodelling of hand representation in adult cortex determined by timing of tactile stimulation. *Nature, 378*(6552), 71–75.

Weinberger, N. M., Javid, R., & Lepan, B. (1995). Heterosynaptic long-term facilitation of sensory-evoked responses in the auditory cortex by stimulation of the magnocellular medial geniculate body in guinea pigs. *Behavioral Neuroscience, 109*(1), 10–17.

Weinshall, D., & Liu, Z. (2000). Mechanisms of generalization in perceptual learning. *Vision Research, 40*(1), 97–109.

Wertheimer, M. (1923). Untersunchungen zur Lehre der Gestalt. *Psychologische Forschung, 4,* 301–350.

Wheeler, D. D. (1970). Processes in the visual recognition of words (Doctoral dissertation, University of Michigan, 1970). *Dissertation Abstracts International, 31*(2), 940B.

Wise, J. A., Kubose, T., Chang, N., Russell, A., & Kellman, P. J. (2000). Perceptual learning modules in mathematics and science instruction. In D. Lemke (Ed.), *Proceedings of the TechEd 2000 Conference,* Amsterdam: IOS Press.