

PERCEPTUAL LEARNING

Robert L. Goldstone

Psychology Building, Indiana University, Bloomington, Indiana 47405;
e-mail: rgoldstone@indiana.edu

KEY WORDS: perception, cognition, discrimination, training, expertise

ABSTRACT

Perceptual learning involves relatively long-lasting changes to an organism's perceptual system that improve its ability to respond to its environment. Four mechanisms of perceptual learning are discussed: attention weighting, imprinting, differentiation, and unitization. By attention weighting, perception becomes adapted to tasks and environments by increasing the attention paid to important dimensions and features. By imprinting, receptors are developed that are specialized for stimuli or parts of stimuli. By differentiation, stimuli that were once indistinguishable become psychologically separated. By unitization, tasks that originally required detection of several parts are accomplished by detecting a single constructed unit representing a complex configuration. Research from cognitive psychology, psychophysics, neuroscience, expert/novice differences, development, computer science, and cross-cultural differences is described that relates to these mechanisms. The locus, limits, and applications of perceptual learning are also discussed.

CONTENTS

INTRODUCTION	586
MECHANISMS OF PERCEPTUAL LEARNING	588
<i>Attentional Weighting</i>	588
<i>Stimulus Imprinting</i>	591
<i>Differentiation</i>	596
<i>Unitization</i>	602
THE LIMITATIONS AND POTENTIAL OF PERCEPTUAL LEARNING	604

INTRODUCTION

The field of perceptual learning has changed significantly since the last *Annual Review of Psychology* chapter—Eleanor Gibson's 1963 review entitled "Perceptual Learning" [reprinted and reappraised by Gibson (1991)]—appeared. Eleanor and James Gibson's ecological approach to perception, with its emphasis on the direct perception of information from the world, has had a profound influence on the direction of the entire field. By this approach, perceptual learning consists of extracting previously unused information (Gibson & Gibson 1955). Identifying what external properties are available to be picked up by people is one of the major research goals. The ecological approach to perceptual learning continues to offer a fertile research program in developmental psychology (Pick 1992) and event perception (Bingham et al 1995, Reed 1996). The focus of the current review will be quite different from a direct perception perspective. The research reviewed here will be predominantly concerned with the internal mechanisms that drive perceptual learning and mediate between the external world and cognition. The bulk of this review will be organized around proposals for specific mechanisms of perceptual adaptation.

Perceptual learning involves relatively long-lasting changes to an organism's perceptual system that improve its ability to respond to its environment and are caused by this environment. As perceptual changes become more ephemeral, the inclination is to speak of adaptation (Helson 1948), attentional processes (Nosofsky 1986), or strategy shifts rather than perceptual learning. If the changes are not due to environmental inputs, then maturation rather than learning is implicated. Perceptual learning may occasionally result in worse performance in perceptual tasks, as is the case with Samuel's (1981) finding that experience with spoken words hinders subjects' decisions about whether they heard white noise or noise combined with speech sounds. Even in this case, experience with words probably increases people's ability to decipher noisy speech, the task with which they are most often confronted. The premise of this definition is that perceptual learning benefits an organism by tailoring the processes that gather information to the organism's uses of the information.

One of the theoretical and empirical challenges underlying the above definition is to distinguish between perceptual and higher-level, cognitive learning. In fact, Hall (1991) has persuasively argued that many results that have been explained in terms of perceptual learning are more parsimoniously described by changes involving the strengthening and weakening of associations. Several strategies have been proposed for identifying perceptual, rather than higher-level, changes. Under the assumption that perception involves the early stages of information processing, one can look for evidence that experience influences early processes. For example, subjective experience not only

alters the perceived colors of familiar objects (Goldstone 1995) but apparently also exerts an influence on color perception before the perceptual stage that creates color after-images has completed its processing (Moscovici & Personnaz 1991). Likewise, experience with the silhouettes of familiar objects exerts an influence before figure/ground segregation is completed (Peterson & Gibson 1994). A second approach to identifying perceptual changes is to observe the time course of the use of particular types of information. For example, on the basis of priming evidence, Sekuler et al (1994) argue that knowledge about what an occluded object would look like if it were completed influences processing after as little as 150 ms. This influence is sufficiently early to typically be counted as perceptual processing. Neurological evidence can provide convergent support to timing studies. For example, practice in discriminating small motions in different directions significantly alters electrical brain potentials that occur within 100 ms of the stimulus onset (Fahle & Morgan 1996). These electrical changes are centered over the primary visual cortex, suggesting plasticity in early visual processing. Karni & Sagi (1991) find evidence, based on the specificity of training to eye (interocular transfer does not occur) and retinal location, that is consistent with early, primary visual cortex adaptation in simple discrimination tasks. Similarly, classical conditioning leads to shifts of neuronal receptive fields in primary auditory cortex toward the frequency of the rewarded tone (Weinberger 1993). In fact, training in a selective attention task may produce differential responses as early as the cochlea, the neural structure that is connected directly to the eardrum via three small bones (Hillyard & Kutas 1983). In short, there is an impressive amount of converging evidence that experimental training leads to changes to very early stages of information processing.

Of the many interesting questions regarding perceptual learning (“What is learned?,” “How long does learning take and last?,” and “How widely does learning transfer?”), this review is organized around “How does learning occur?” Consequently, a wide range of fields that investigate mechanisms underlying perceptual learning will be surveyed. Evidence from developmental psychology is very important because many of the most dramatic changes to human perceptual systems occur within the first seven years of life (Aslin & Smith 1988). Neuroscience provides concrete mechanisms of adaptation, and the field of neural plasticity has recently experienced tremendous growth (Kolb 1995, McGaugh et al 1995). Analyses of expertise and cross-cultural comparisons assess the perceptual impact of extended environmental influences. Researchers in computer science have made valuable contributions to our understanding of human psychology by describing functional algorithms for adaptation in networks involving many interacting units. In many cases, perceptual changes that have been empirically observed through studies of ex-

perts, laboratory training studies, and different cultures are given concrete accounts by computational and neural models.

MECHANISMS OF PERCEPTUAL LEARNING

Perceptual learning is not achieved by a unitary process. Psychophysicists have distinguished between relatively peripheral, specific adaptations and more general, strategic ones (Doane et al 1996, Sagi & Tanne 1994), and between quick and slow perceptual learning processes (Karni & Sagi 1993). Cognitive scientists have distinguished between training mechanisms driven by feedback (supervised training) and those that require no feedback, instead operating on the statistical structure inherent in the environmentally supplied stimuli (unsupervised training). Organizing perceptual learning in terms of mechanisms rather than domains results in some odd couplings (linking, for example, neuroscientific and cross-cultural studies bearing on perceptual differentiation), but has the advantage of connecting phenomena that are deeply related and may inform each other.

Attentional Weighting

One way in which perception becomes adapted to tasks and environments is by increasing the attention paid to perceptual dimensions and features that are important, and/or by decreasing attention to irrelevant dimensions and features. A feature is a unitary stimulus element, whereas a dimension is a set of linearly ordered features. “3 centimeters” and “red” are features; length and color are dimensions.

Attention can be selectively directed toward important stimulus aspects at several different stages in information processing. Researchers in animal learning and human categorization have described shifts toward the use of dimensions that are useful for tasks (Nosofsky 1986) or have previously been useful (Lawrence 1949). Lawrence describes these situations as examples of stimulus dimensions “acquiring distinctiveness” if they have been diagnostic in predicting rewards. Nosofsky describes attention shifts in terms of psychologically “stretching” dimensions that are relevant for categorizations. During category learning, people show a trend toward deemphasizing preexperimentally salient features, and emphasizing features that reliably predict experimental categories (Livingston & Andrews 1995). The stimulus aspects that are selectively attended may be quite complex; even pigeons can learn to selectively attend to the feature “contains human” in photographs (Herrnstein 1990). In addition to important dimensions acquiring distinctiveness, irrelevant dimensions also acquire equivalence, becoming less distinguishable (Honey & Hall 1989). For example, in a phenomenon called “latent inhibi-

tion,” stimuli that are originally varied independently of reward are harder to later associate with reward than those that are not initially presented at all (Lubow & Kaplan 1997, Pearce 1987). Haider & Frensch (1996) find that improvements in performance are frequently due to reduced processing of irrelevant dimensions.

The above studies illustrate shifts in the use of dimensions as a function of their task relevance, but these shifts may be strategic choices rather than perceptual in nature. One source of evidence that they are not completely voluntary is that attentional highlighting of information occurs even if it is to the detriment of the observer. When a letter consistently serves as the target in a detection task and then later becomes a distractor—a stimulus to be ignored—it still automatically captures attention (Shiffrin & Schneider 1977). The converse of this effect, negative priming, also occurs: targets that were once distractors are responded to more slowly than never-before-seen items (Tipper 1992). In the negative priming paradigm, the effect of previous exposures of an item can last upward of two weeks (Fox 1995), suggesting that a relatively permanent change has taken place.

CATEGORICAL PERCEPTION A phenomenon of particular interest for attentional accounts of perceptual adaptation is categorical perception. According to this phenomenon, people are better able to distinguish between physically different stimuli when the stimuli come from different categories than when they come from the same category (Calder et al 1996; see Harnad 1987 for several reviews of research). The effect has been best documented for speech phoneme categories. For example, Liberman et al (1957) generated a continuum of equally spaced consonant-vowel syllables going from /be/ to /de/. Observers listened to three sounds—A followed by B followed by X—and indicated whether X was identical to A or B. Subjects performed the task more accurately when syllables A and B belonged to different phonemic categories than when they were variants of the same phoneme, even when physical differences were equated.

There is evidence that some categorical perception effects are not learned but are either innate or a property of the acoustical signal itself. Infants as young as 4 months showed categorical perception for speech sounds (Eimas et al 1971), and even chinchillas (Kuhl & Miller 1978) and crickets (Wytenbach et al 1996) show categorical perception effects for sound.

Still, recent evidence has indicated that sound categories, and categorical perception more generally, are subject to learning (Lively et al 1993). Whether categorical perception effects are found at particular physical boundaries depends on the listener's native language. In general, a sound difference that crosses the boundary between phonemes in a language will be more dis-

criminally to speakers of that language than to speakers of a language in which the sound difference does not cross a phonemic boundary (Repp & Liberman 1987, Strange & Jenkins 1978). Laboratory training on the sound categories of a language can produce categorical perception among speakers of a language that does not have these categories (Pisoni et al 1982). Expert musicians, but not novices, show a pronounced categorical perception effect for relative pitch differences, suggesting that training was instrumental in sensitizing boundaries between semitones (Burns & Ward 1978, Zatorre & Halpern 1979). A visual analog exists: faces for which subjects are “experts,” familiar faces, show categorical perception (increased sensitivity to differences at the half-way point between the faces) as one familiar face is transformed into another familiar face; however, no categorical perception is found for unfamiliar faces (Beale & Keil 1995).

There are several ways that physical differences between categories might become emphasized relative to within-category differences. In support of the possibility that people lose their ability to make within-category discriminations, very young infants (2 months old) show sensitivity to differences between speech sounds that they lose by the age of 10 months (Werker & Lalonde 1988, Werker & Tees 1984). This desensitization only occurs if the different sounds come from the same phonetic category of their native language. However, given the difficulty in explicitly instructing infants to respond to physical rather than phonetic differences between sounds, these results should be conservatively interpreted as showing that physical differences that do not make a functional difference to children become perceptually or judgmentally deemphasized. Laboratory experiments by Goldstone (1994) have suggested that physical differences between categories become emphasized with training. After learning a categorization in which one dimension was relevant and a second dimension was irrelevant, subjects were transferred to same/different judgments (“Are these two squares physically identical?”). Ability to discriminate between stimuli in the same/different judgment task was greater when they varied along dimensions that were relevant during categorization training, and was particularly elevated at the boundary between the categories. Further research showed that category learning systematically distorts the perception of category members by shifting their perceived dimension values away from members of opposing categories (Goldstone 1995). In sum, there is evidence for three influences of categories on perception: (a) category-relevant dimensions are sensitized, (b) irrelevant variation is deemphasized, and (c) relevant dimensions are selectively sensitized at the category boundary.

Computational efforts at explaining categorical perception have mainly centered on neural networks. In two such models, equally spaced stimuli along a continuum are associated with category labels, and the networks adapt their

input-to-category connections so that the stimuli come to evoke their correct category assignment (Anderson et al 1977, Harnad et al 1995). In effect, the category feedback establishes attractor states that pull the different members of a category to a common point, thereby reducing their distinctiveness.

Stimulus Imprinting

A second way that perception can adapt to an environment is by directly imprinting to it. Through imprinting, detectors (also called receptors) are developed that are specialized for stimuli or parts of stimuli. The term imprinting captures the idea that the form of the detector is shaped by the impinging stimulus. Internalized detectors develop for repeated stimuli, and these detectors increase the speed, accuracy, and general fluency with which the stimuli are processed. Although evidence for neural implementations of acquired detectors will be considered, more generally the reviewed studies support functional detectors—any abstract device or process that explains the selective benefit to important, repeated patterns.

WHOLE STIMULUS STORAGE Imprinting may occur for entire stimuli, in which case a receptor develops that internalizes specific instances. Models that preserve stimuli in their entirety are called exemplar (Nosofsky 1986) or instance-based (Logan 1988) models. For example, in Logan's model, every exposure to a stimulus leads to an internalized trace of that stimulus. As more instances are stored, performance improves because more relevant instances can be retrieved, and the time required to retrieve them decreases. Instance-based models are supported by results showing that people's performance in perceptual tasks is closely tied to their amount of experience with a particular stimulus. Consistent with this claim, people can identify spoken words more accurately when they are spoken by familiar voices (Palmeri et al 1993). Doctors' diagnoses of skin disorders are facilitated when they are similar to previously presented cases, even when the similarity is based on attributes that are irrelevant for the diagnosis (Brooks et al 1991). Increasing the frequency of a cartoon face in an experiment increases its classification accuracy (Nosofsky 1991). After several hours of training in a numerosity judgment task ("How many dots are there?"), people's response times are the same for all levels of numerosity between 6 and 11 dots, but only for dots that are arranged as they were during training (Palmeri 1997), consistent with the notion that slow counting processes can be circumvented by storing specific arrangements of dots. Even when people know a simple, clear-cut rule for a perceptual classification, performance is better on frequently presented items than rare items (Allen & Brooks 1991). Thus, even in situations where one might think abstract or rule-based processes are used, there is good evidence

that observers become tuned to the particular instances to which they are exposed.

People are better able to perceptually identify unclear or quickly presented stimuli when they have been previously exposed to them. Although this effect is traditionally discussed in terms of implicit memory for exposed items, it also provides a robust example of perceptual learning. The identification advantage for familiarized instances lasts at least three weeks, requires as few as one previous presentation of an item, and is often tied to the specific physical properties of the initial exposure of the item (Schacter 1987). In brief, instance memories that are strong and quickly developed facilitate subsequent perceptual tasks involving highly similar items.

The power of instance-based models has not been ignored by object recognition researchers. This has led to a renewed interest in the recently dismissed class of “template” models. According to these models, objects are recognized by comparing them to stored, photograph-like images (templates) of known objects. Objects are placed into the same category as the template to which they are most similar. In some cases, preprocessing operations rotate and distort templates to maximize their overlap with the presented object (Hinton et al 1992). Ullman (1989) has shown that template models can be highly effective, and that preprocessing operations can find good matches between an object and template without knowing ahead of time what the object is, as long as at least three points of alignment between the object and template can be found on the basis of local physical cues. Poggio & Edelman (1990) present a neural network model that learns to recognize three-dimensional objects by developing units specialized for presented two-dimensional views, associating them with their correct three-dimensional interpretation, and interpolating between stored views for recognizing novel objects. Consistent with this model’s assumption that receptors become tuned to particular viewpoints, humans can learn to identify three-dimensional objects by seeing two-dimensional views that have been arbitrarily paired with the three-dimensional object (Sinha & Poggio 1996). Tarr (1995) provides support for the storage of multiple views to aid recognition by showing that the time to recognize rotated objects is a function of their rotational distance to the nearest stored viewpoint.

FEATURE IMPRINTING Rather than imprinting on entire stimuli, there is also evidence that people imprint on parts or features of a stimulus. If a stimulus part is important, varies independently of other parts, or occurs frequently, people may develop a specialized detector for that part. This is a valuable process because it leads to the development of new “building blocks” for describing stimuli (Schyns et al 1998, Schyns & Murphy 1994). Parts that are devel-

oped in one context can be used to efficiently describe subsequent objects. Efficient representations are promoted because the parts have been extracted because of their prevalence in an environment, and thus are tailored to the environment.

Schyns & Rodet (1997) find that unfamiliar parts (arbitrary curved shapes within an object) that are important in one task are more likely to be used to represent subsequent categories. Their subjects were more likely to represent a conjunction of two parts, X and Y, in terms of these two components (rather than as a whole unit, or a unit broken down into different parts) when they received previous experience with X as a defining part for a different category. Configurations of dots are more likely to be circled as coherent components of patterns if they were previously important for a categorization (Hock et al 1987). Likewise, Hofstadter (1995) and his colleagues describe how learning to interpret an object as possessing certain parts creates a bias to see other objects in terms of those parts.

Several computational models have been recently devised that create perceptual building blocks during the course of being exposed to, or categorizing, objects. Neural networks have been particularly popular because they often possess hidden units that intervene between inputs and outputs and can be interpreted as developing internal representations of presented inputs (Rumelhart et al 1986). These internal representations can function as acquired feature detectors, built up through environmental exposure. For example, simple exposure to photographs of natural scenes suffices to allow neural networks to create a repertoire of oriented line segments to be used to describe the scenes (Miikkulainen et al 1997, Schmidhuber et al 1996). These feature detectors bear a strong resemblance to neural detectors found in the primary visual cortex and are created by learning algorithms that develop units that respond to independent sources of regularity across photographs. Networks with detectors that adapt by competing for the privilege to accommodate inputs can generate specialized detectors resembling ocular dominance and orientation columns found in the visual cortex (Obermayer et al 1995). These networks do not require feedback labels or categorizations; the stimuli themselves contain sufficient regularities and redundancies that can be exploited to generate efficient vocabularies (Grossberg 1991). However, if neural networks do receive feedback about stimulus categorizations, then the features that they develop can be tailored to these categories (Intrator 1994, Rumelhart et al 1986). The simplicity, predictive power, and value of neural networks that create their own featural descriptions make these systems exciting and fruitful avenues for exploration.

There is also neurological evidence for perceptual learning via imprinting on specific features within a stimulus. Weinberger (1993) reviews evidence

that cells in the auditory cortex become tuned to the frequency of often-repeated tones. Ascending in complexity, cells in the inferior temporal cortex can be tuned by extended experience (about 600,000 trials) with 3D objects (Logothetis et al 1995); these cells also show heightened response to novel views of the trained object. Cells in this same area can be highly selective for particular faces, and this specificity is at least partially acquired given that it is especially pronounced for familiar faces (Perrett et al 1984).

The cognitive, computational, and neurophysiological results indicate that the “building blocks” used to describe objects are adapted to environmental inputs. In many of the cases considered thus far, feature and part detectors are devised that capture the regularities implicit in the set of input stimuli. However, the detectors that develop are also influenced by task requirements and strategies. For example, altering the color of target objects from training to transfer does not influence performance unless the training task requires encoding of color (Logan et al 1996). In general, whether a functional detector is developed will depend on both the objective frequency and subjective importance of the physical feature (Sagi & Tanne 1994, Shiu & Pashler 1992). Systems that can acquire new feature detectors have functional advantages over systems that employ a hard-wired set of detectors. One difficulty with fixed sets of features is that it is hard to choose exactly the right set of elements that will suffice to accommodate all possible future entities. On the one hand, if a small set of primitive elements is chosen, then it is likely that two entities will eventually arise that must be distinguished, but cannot with any combination of available primitives. On the other hand, if a set of primitives is sufficiently large to construct all entities that might occur, then it will likely include many elements that lie unused, waiting for their moment of need to possibly arise (Schyns et al 1998). However, by developing new elements as needed, newly important discriminations can cause the construction of detectors that are tailored for the discrimination.

TOPOLOGICAL IMPRINTING A third type of imprinting occurs at a more abstract level. Rather than developing detectors for particular stimuli or features, environmental regularities that span across a set of stimuli can also be internalized. The patterns impinging upon an organism will have certain similarities to one another. These similarities can be represented by plotting each pattern in a multidimensional space. Topological imprinting occurs when the space and the positions of patterns within the space are learned as a result of training with patterns. Rather than simply developing independent detectors, topological imprinting implies that a spatially organized network of detectors is created.

The simplest form of topological imprinting is to create a set of feature values ordered along a single dimension. Developmental evidence suggests that dimensional organizations are learned. On the basis of evidence from a "Which is more?" task, children and adults agree that large objects are "more" than small objects, but three-year-old children treat dark objects as more than light objects, unlike adults (Smith & Sera 1992). Loudness is originally disorganized for children, but comes to be dimensionally organized with loud sounds being perceived as more than soft sounds. The importance of dimensionally organized percepts is apparent from Bedford's (1993, 1995) work on learning the relations between dimensions. She argues that perceptual learning involves adaptively mapping from one dimension to another. For example, upon wearing prism eyeglasses that distort the relation between visual information and proprioceptive feedback, learning is much easier when the entire visual dimension can be shifted or warped to map onto the proprioceptive dimension than when unrelated visual-motor associations must be acquired. Both experiments point to people's natural tendency to draw associations between dimensions. One of the most striking examples of this phenomenon continues to be Howells' (1944) experiment in which people learn to associate a particular tone with the color red after several thousand trials, and then are transferred to a task where they try to identify a neutral color white. When the tone is present, people systematically choose as white a color that is slightly green, suggesting that the tone has come to substitute for redness to some extent. Perceptual learning involves developing dimensional structures and also mappings across these dimensions.

Quite a bit is known about the neural and computational mechanisms underlying the acquisition of topologically structured representations of the environment. Sensory maps in the cortex preserve topological structures of the peripheral sensory system; for example, the primary sensory area responsible for the middle finger (digit 3) of the Macaque monkey lies between the areas responsible for digits 2 and 4. Several types of adaptive cortical change, all of which preserve topological mapping, are observed when environmental or cortical changes occur (Garraghty & Kaas 1992). When cortical areas are lesioned, neighboring areas newly respond to sensory information formerly controlled by the lesioned area; when external sensory organs are disabled, cortical areas formerly activated by the organ become sensitive to sensory stimulation formerly controlled by its neighboring areas (Kaas 1991). When two fingers are surgically fused, creating highly correlated inputs, a large number of cortical areas develop that respond to both fingers (Allard et al 1991). Kohonen (1995) has developed a framework for describing neural networks that develop topological structures with learning. These networks are composed of detectors that compete for the opportunity to learn to respond to inputs more

strongly, and are arranged in topologies (typically, two-dimensional lattices). These topologies influence learning—not only does the unit that is best adapted to an input learn to respond more vigorously to the input, but so do its neighbors. Variants of Kohonen’s networks can acquire topologies similar to those found in the cortex, and can adapt in similar ways to network lesions and alterations in the environment (Miikkulainen et al 1997). Other neural networks capture more abstract spatial dimensions, learning dimensions that optimally describe the similarities between a set of objects (Edelman & Intrator 1997). In general, these networks develop detectors that are locally tailored to particular inputs, and also arrange their detectors in a global configuration that represents similarities and dimensions across inputs.

Differentiation

A major mechanism of perceptual learning is for percepts to become increasingly differentiated from each other. By differentiation, stimuli that were once psychologically fused together become separated. Once separated, discriminations can be made between percepts that were originally indistinguishable. As with imprinting, differentiation occurs at the level of whole stimuli and features within stimuli.

DIFFERENTIATION OF WHOLE STIMULI In the classic examples of wine experts learning to distinguish the upper and lower halves of a bottle of Madeira by taste, poultry sorters learning to distinguish male from female chicks, and parents learning to uniquely identify their identical twin children, perceptual adaptation involves developing increasingly differentiated object representations. In many cases, simple preexposure to the stimuli to be distinguished promotes their differentiation. Rats who have cutout shapes visible from their cages are better able to learn subsequent discriminations involving these shapes than rats who are exposed to other shapes (Gibson & Walk 1956). Practice in identifying visual “scribbles” increases their discriminability, even when no feedback is provided (Gibson & Gibson 1955). However, learning to differentiate between objects is typically accelerated by training in which the objects are associated with different labels or responses (Gibson 1969, Hall 1991).

Psychophysical differentiation Laboratory studies have extensively studied training effects involving simple discriminations. In vernier discrimination tasks, subjects respond whether one line is displaced above or below a second line. Training in this task can produce impressive improvements, to the point that subjects exhibit resolution finer than the spacing between individual pho-

toreceptors (Poggio et al 1992). Such hyperacuity is possible because receptive fields of cells overlap considerably, and thus points that fall within the receptive field of one cell can be discriminated by differential impacts on other cells. Discrimination training is often highly specific to the task. Trained performance on a horizontal discrimination task frequently does not transfer to a vertical version of the same task (Fahle & Edelman 1993, Poggio et al 1992), does not transfer to new retinal locations (Fahle et al 1995, Shiu & Pashler 1992), and does not completely transfer from the trained eye to the untrained eye (Fahle et al 1995).

The surprising specificity of simple discrimination learning has led some researchers to posit an early cortical locus of adaptation, perhaps as early as the primary visual cortex (Gilbert 1996, Karni & Sagi 1991). Improvement in the discrimination of motion of a random dot field has been shown to be associated with a change in the response characteristics of individual cells in area MT in the parietal cortex (Zohary et al 1994). Computational models have explained improvements in discrimination training in terms of changes in weights between cells and output units that control judgments (Poggio et al 1992). Each cell has a limited receptive field and specific orientation, and cells that predict vernier discriminations become more influential over time. Thus, the proposed mechanism for differentiation is selective emphasis of discriminating receptive cells.

A related method for implementing differentiation is to develop expanded representations for receptive cells that permit discrimination of objects that should receive different responses. Monkeys trained to make discriminations between slightly different sound frequencies develop larger cortical representations for the presented frequencies than control monkeys (Recanzone et al 1993). Similarly, monkeys learning to make a tactile discrimination with one hand develop a larger cortical representation for that hand than for the other hand (Recanzone et al 1992). Elbert et al (1995) measured brain activity in the somatosensory cortex of violinists as their fingers are lightly touched. There was greater activity in the sensory cortex for the left hand than the right hand, consistent with the observation that violinists require fine movements of their left-hand fingers considerably more than their right-hand fingers.

A third neural mechanism for stimulus differentiation is to narrow the tuning of critical receptors. Receptors that are originally broadly tuned (large receptive fields) often become responsive to an increasingly limited range of stimuli with training. Recanzone et al (1993) observe a narrowing of frequency-sensitive receptors following auditory discrimination. Saarinen & Levi (1995) also find evidence that training in a vernier discrimination task results in receptors that are more narrowly tuned to diagnostic orientations. A

mechanism for differentiation explored by Luce et al (1976) is that a roving attentional band can be selectively placed on critical regions of a perceptual dimension, and that signals falling within the band are given a sensory representation about an order of magnitude greater than signals falling outside of the band. These four mechanisms—selective weighting of discriminating cells, expanding regions dedicated to discriminating cells, narrowing tuning of discriminating cells, and shifting attentional “magnifiers” to critical regions—all serve to differentiate stimuli by psychologically warping local regions of stimulus space.

Differentiation of complex stimuli Differentiation of more complex stimuli that differ across many dimensions has also been studied. Lively et al (1993) report training procedures that allow Japanese speakers to acquire a discrimination between the phonemes /r/ and /l/ that is not present in their native language. The methodological innovations apparently needed to assure general transfer performance are to provide the phonemes in natural words, to give listeners words spoken by many different speakers, and to give immediate feedback as to the correct word. A general finding has been that increasing the variability of instances within the categories to be discriminated increases the amount of training time needed to reach a criterion level of accuracy, but also yields better transfer to novel stimuli (Posner & Keele 1968). Another result of practical interest is that discrimination performance can be improved by an “easy-to-hard” procedure in which subjects are first exposed to easy, highly separated discriminations along a dimension (such as black versus white stimuli on the dimension of brightness), and then are given successively more difficult discriminations along the same dimension (Mackintosh 1974). Apparently, first presenting the easy discrimination allows organisms to allocate attention to the relevant dimension.

A major subfield within stimulus differentiation has explored expertise in face perception. People are better able to identify faces belonging to races with which they are familiar (Shapiro & Penrod 1986). For example, in general, Caucasian participants in the United States are better able to identify Caucasian faces than African-American faces. This is another example of familiar objects becoming increasingly differentiated. A common account for the difficulty in recognizing cross-race faces is that people become adept at detecting the features that are most useful in distinguishing among faces we commonly see (O’Toole et al 1996). Interestingly, people are faster at categorizing those faces that are more difficult to identify. For example, in an African-American/Caucasian discrimination task, Caucasian participants are faster at categorizing African-Americans (as African-Americans) than Caucasians (Valentine 1991). Valentine explains this effect in terms of greater perceived distances

between familiar faces, which slows tasks such as a two-category discrimination that require treating familiar faces as equivalent. In contrast, Levin (1996) obtains evidence that African-American categorizations are facilitated for Caucasians because of a quickly coded race feature that marks cross-race but not same-race faces. This later possibility suggests that object differentiation may be achieved by developing features that uniquely pick out less common objects from familiar objects (Goldstone 1996), and is consistent with results showing that perceptual retention of abnormal chest X-rays increases with radiological expertise, whereas retention of normal X-rays actually decreases with expertise (Myles-Worsley et al 1988). Levin's account is not necessarily incompatible with the standard account; features may become salient if they serve to either discriminate among familiar objects or to distinguish rare objects from familiar ones.

The results are somewhat mixed with respect to the benefit of instructional mediation in learning to differentiate stimuli. For learning to discriminate between beers (Peron & Allen 1988), experience with tasting beers improved performance, but increased experience with beer-flavor terminology did not. However, in learning to sort day-old chickens by gender, college students with no experience were able to categorize photographs of chickens nearly as well as were expert chicken sorters if they were given a short page of instructions describing shape-based differences between male and female chickens (Biederman & Shiffrar 1987). It is an open question whether genuinely perceptual changes can be produced after simply reading a brief set of instructions. Those who argue that perceptual phenomena are generally not highly malleable to instructions and strategies (Rock 1985) might consider Biederman & Shiffrar's results to merely show that perceptual features that have previously been learned can become linked to categories by instructions. On the other hand, strategic intentions and labels can produce phenomenologically different percepts of ambiguous objects, and the linguistic labels chosen to describe an object can radically reorganize its perception (Wisniewski & Medin 1994). The possibility that perceptual processes are altered by instructional or strategic manipulations cannot be dismissed.

Differentiation of categories Ascending even further in terms of the complexity of stimuli to be differentiated, not only do simple and complex objects become differentiated with experience, but entire categories do as well. Category learning often involves dividing a large, general category into subcategories. Intuition tells us that experts in a field have several differentiated categories where the novice has only a single category. Empirical support for this notion comes from Tanaka & Taylor's (1991) study of speeded classification by dog and bird experts. Categories can be ordered in terms of specificity, from

highly general superordinate categories (e.g. “animal”), to moderately specific basic-level categories (“dog”), to highly specific subordinate categories (“German shepherd”). When shown photographs of objects and asked to verify whether they belong to a particular category, experts are able to categorize at basic and subordinate levels equally quickly, but only for the objects within their domain of expertise. In contrast, novices (e.g. bird experts shown dog photographs) show a pronounced advantage for basic-level categorizations. Extending the previously described identification advantage for same-race faces, O’Toole et al (1996) find that Caucasians and Japanese are faster at classifying faces of their own race into “male” and “female” categories than faces of the other race. Categories, not just objects, are more differentiated within familiar domains.

Cross-cultural differences provide additional evidence that categories that are important become highly differentiated (Malt 1995). For example, the Tzeltal Indians group all butterflies together in one category, but have 16 different categories for their larvae, which are culturally important as food sources and crop pests (Hunn 1982). The observer-relative terms “Left” and “Right” are important spatial concepts in some cultures, whereas other cultures (e.g. speakers of Guugu Yimithrr) much more frequently describe space in terms of absolute directions such as “North” and “South.” Levinson (1996) argues that this cultural difference has an influence on perceptual tasks such as completing paths and discriminating between events that differ as to their relative or absolute spatial relations. Generally, the degree of differentiation among the categories of a culture is a joint function of the importance of the categories for the culture and the objective number and frequency of the categories in the environment (Geoghegan 1976).

There is also developmental evidence that categories become more differentiated with age. Infants tend to successively touch objects that are perceptually similar. Using successive touching as an indicator of subjective groupings, Mandler et al (1991) show that 18-month-old infants group objects at the superordinate level (e.g. successively touching toy goats and cats more frequently than dogs and planes) before they show evidence of basic-level categories (e.g. by successively touching two cats more frequently than a cat and a dog). In sum, evidence from expert/novice differences, cross-cultural differences, development, and neuroscience (Farah 1990) provide converging evidence that broader levels of categorization are deeply entrenched and perhaps primary, and that experience yields more subtly differentiated categories.

DIFFERENTIATION OF DIMENSIONS Just as experience can lead to the psychological separation of stimuli or categories, it can also lead to the separation of

perceptual dimensions that comprise a single stimulus. Dimensions that are originally treated as fused often become segregated with development or training. People often shift from perceiving stimuli in terms of holistic, overall aspects to analytically decomposing objects into separate dimensions.

This trend has received substantial support from developmental psychology. Evidence suggests that dimensions that are easily isolated by adults, such as the brightness and size of a square, are treated as fused together for children (Smith 1989a). It is relatively difficult for young children to say whether two objects are identical on a particular property, but relatively easy for them to say whether they are similar across many dimensions (Smith 1989a). Children have difficulty identifying whether two objects differ on their brightness or size even though they can easily see that they differ in some way (Kemler 1983). Children also show considerable difficulty in tasks that require selective attention to one dimension while ignoring another (Smith & Evans 1989). When given the choice of sorting objects by their overall similarity or by selecting a single criterial dimension, children tend to use overall similarity, whereas adults use the single dimension (Smith 1989b). Perceptual dimensions seem to be more tightly integrated for children than adults, such that children cannot easily access the individual dimensions that compose an object.

The developmental trend toward differentiated dimensions is echoed by adult training studies. In certain circumstances, color experts (art students and vision scientists) are better able to selectively attend to dimensions (e.g. hue, chroma, and value) that comprise color than are nonexperts (Burns & Shepp 1988). People who learn a categorization in which color saturation is relevant and color brightness is irrelevant develop selectively heightened sensitivity at making saturation discriminations (Goldstone 1994), even though prior to training it is difficult for adults to selectively attend to brightness without attending to saturation. Melcher & Schooler (1996) provide suggestive evidence that expert, but not nonexpert, wine tasters isolate independent perceptual features in wines that closely correspond to the terminology used to describe wines.

Several computational models have been proposed for differentiation. Competitive learning networks differentiate inputs into categories by specializing detectors to respond to classes of inputs. Random detectors that are slightly more similar to an input than other detectors will learn to adapt themselves toward the input and will inhibit other detectors from doing so (Rumelhart & Zipser 1985). The end result is that originally similar detectors that respond almost equally to all inputs become increasingly specialized and differentiated over training. Detectors develop that respond selectively to particular classes of input patterns or dimensions within the input. Smith et al (1997) present a

neural network simulation of the development of separated dimensions in children. In the network, dimensions become separated by detectors developing strong connections to specific dimensions while weakening their connections to all other dimensions. The model captures the empirical phenomenon that dimension differentiation is greatly facilitated by providing comparisons of the sort “this red square and this red triangle have the same color.”

Unitization

Unitization is a perceptual learning mechanism that seemingly operates in a direction opposite to differentiation. Unitization involves the construction of single functional units that can be triggered when a complex configuration arises. Via unitization, a task that originally required detection of several parts can be accomplished by detecting a single unit. Whereas differentiation divides wholes into cleanly separated parts, unitization integrates parts into single wholes.

In exploring unitization, LaBerge (1973) found that when stimuli were unexpected, participants were faster at responding to actual letters than to letter-like controls. Furthermore, this difference was attenuated as the unfamiliar letter-like stimuli became more familiar with practice. He argued that the components of often-presented stimuli become processed as a single functional unit when they consistently occur together. More recently, Czerwinski et al (1992) have described a process in which conjunctions of stimulus features are “chunked” together so that they become perceived as a single unit. Shiffrin & Lightfoot (1997) argued that even separated line segments can become unitized following prolonged practice with the materials. Their evidence comes from the slopes relating the number of distractor elements to response time in a feature search task. When participants learned a conjunctive search task in which three line segments were needed to distinguish the target from distractors, impressive and prolonged decreases in search slopes were observed over 20 sessions.

Other evidence for unitization comes from word perception. Researchers have argued that words are perceived as single units because of people’s life-long experience with them. These word units can be processed automatically and interfere with other processes less than do nonwords (O’Hara 1980, Smith & Haviland 1972). Results have shown that the advantages attributable to words over nonwords cannot be explained by the greater informational redundancy of letters within words (Smith & Haviland 1972). Instead, these researchers argue for recognition processes that respond to information at levels higher than the individual letters. Salasoo et al (1985) find that the advantage of words over nonwords in perceptual identification tasks can be eliminated by

repetitively exposing participants to the stimuli. They explain their results in terms of developing single, unitized codes for repeated nonwords.

Evidence for unitization also comes from researchers exploring configural perception. For example, researchers have argued that faces are processed in a holistic or configural manner that does not involve analyzing faces into specific features (Farah 1992). According to the “inversion effect” in object recognition, the recognition cost of rotating a stimulus 180 degrees in the picture plane is much greater for specialized, highly practiced stimuli than for less specialized stimuli (Diamond & Carey 1986, Tanaka & Gauthier 1997). For example, recognition of faces is substantially less fast and accurate when the faces are inverted. This large difference between upright and inverted recognition efficiency is not found for other objects and is not found to the same degree for less familiar cross-race faces. Diamond & Carey (1986) report a large inversion cost for dog breed recognition, but only for dog experts. Similarly, Gauthier & Tarr (1997) report that large inversion costs for a nonsense object can be created in the laboratory by giving participants prolonged exposure to the object. They conclude that repeated experience with an object leads to developing a configural representation of it that combines all of its parts into a single, viewpoint-specific, functional unit.

There is also evidence that children develop increasingly integrated representations of visual objects as they mature. Whereas three-year-old children tend to break objects into simple, spatially independent parts, five-year-olds use more complicated spatial relations to connect the parts together (Stiles & Tada 1996). It has even been claimed that configural association systems require about 4.5 years to develop, and before this time, children can solve perceptual problems requiring elements but not configurations of elements (Rudy et al 1993).

Computer and neural sciences have provided insights into methods for implementing unitization. Grossberg’s self-organizing ART systems (Grossberg 1984, 1991) create units by building bidirectional links between several perceptual features and a single unit in a deeper layer of the neural network. Triggering the single unit suffices to reproduce the entire pattern of perceptual features. Mozer et al (1992) develop a neural network that creates configural units by synchronizing neurons responsible for visual parts to be bound together. Visual parts that co-occur in a set of patterns will tend to be bound together, consistent with the evidence above indicating that units are created for often-repeated stimuli. Neural mechanisms for developing configural units with experience are located in the superior colliculus and inferior temporal regions. Cells in the superior colliculus of several species receive inputs from many sensory modalities (e.g. visual, auditory, and somatosensory), and differences in their activities reflect learned associations across these modalities (Stein &

Wallace 1996). Within the visual modality, single cells of the inferior temporal cortex become selectively responsive to complex objects that have been repetitively presented (Logothetis et al 1995).

Unitization may seem at odds with dimension differentiation. There is an apparent contradiction between experience creating larger “chunks” via unitization and dividing an object into more clearly delineated parts via differentiation. This incongruity can be transformed into a commonality at a more abstract level. Both mechanisms depend on the requirements established by tasks and stimuli. Objects will tend to be decomposed into their parts if the parts reflect independent sources of variation, or if the parts differ in their relevancy (Schyns & Murphy 1994). Parts will tend to be unitized if they co-occur frequently, with all parts indicating a similar response. Thus, unitization and differentiation are both processes that build appropriate-sized representations for the tasks at hand. Both phenomena could be incorporated in a model that begins with a specific featural description of objects, and creates units for conjunctions of features if the features frequently occur together, and divides features into subfeatures if independent sources of variation within the original features are detected.

THE LIMITATIONS AND POTENTIAL OF PERCEPTUAL LEARNING

Thus far, the reviewed evidence has focused on positive instances of perceptual learning—situations where training produces changes, oftentimes strikingly large, to our perceptual systems. However, a consideration of the limits on perceptual learning leads to a better understanding of the constraints on learning, and hence of the mechanisms that are at work when learning is achieved.

Previously reviewed evidence suggests strong limits on the generality of perceptual learning. Training on simple visual discriminations often does not transfer to different eyes, to different spatial locations, or to different tasks involving the same stimuli (Fahle & Morgan 1996, Shiu & Pashler 1992). As suggested by the strong role played by imprinting, perceptual learning often does not transfer extensively to new stimuli or tasks different from those used during training. Several researchers have argued that generalization between tasks is only found to the extent that the tasks share procedural elements in common (Anderson 1987, Kolers & Roediger 1984). At the same time, perceptual training does often transfer not just within a sensory modality, but across sensory modalities. Training on a visual discrimination involving certain shapes improves performance on a later tactile discrimination involving the same shapes (Hughes et al 1990). Not only does cross-modality transfer occur, but it has also been shown computationally that two modalities that are trained

at the same time and provide feedback for each other can reach a level of performance that would not be possible if they remained independent (Becker 1996, de Sa & Ballard 1997, Edelman 1987). Consistent with these arguments for mutually facilitating modalities, children with auditory deficits but normal IQs also tend to show later deficits in visual selective attention tasks (Quittner et al 1994). One principle for unifying some of the evidence for and against generalization of training seems to be that when perceptual learning involves changes to early perceptual processes, then there will be less generalization of that learning to other tasks (Sagi & Tanne 1994, Sireteanu & Rettenbach 1995).

In addition to constraints on generalization, there are limits on whether perceptual learning occurs at all. To take unitization as an example, many studies indicate a surprising inability of people to build single chunks out of separated dimensions. In Treisman & Gelade's (1980) classic research on feature search, the influence of distractor letters in a conjunctive search remained essentially unchanged over 1664 trials, suggesting that new functional units cannot be formed for conjunctions of color and shape. Although these results are replicable, they depend on the particular features to be joined together. Shiffrin & Lightfoot (1997) report five-fold improvements in response times in a similar, conjunctive search paradigm in which the conjunctions are defined not by color and shape but by different line segments. Searching for conjunctions of shape parts that are formed by life-long experience with letters (Wang et al 1994), or brief laboratory experience (Lubow & Kaplan 1997), is quite different from searching for unfamiliar conjunctions. The influence of distractors on a conjunctive task involving relations such as "dash above plus" is not modulated by practice if the dash and plus are disconnected, but if they are connected, then pronounced practice effects are observed (Logan 1994). People are much more adept at learning conjunctions between shape and position than between shape and color, even when position and color are equally salient (Saiki & Hummel 1996). Thus, logically equivalent conjunctive search tasks can produce very widely different perceptual learning patterns depending on the conjoined features. Features or dimensions that are similar to each other are easy to join together and difficult to isolate (Melara & Marks 1990), and perceptual learning is constrained by these relations.

Perceptual learning at any given time is always constrained by the existing structure of the organism. As such, it is misguided to view perceptual learning as the opposite of innate disposition. Although apparently paradoxical, it is the constraints of systems that allow for their adaptation. Eimas (1994, 1997) provides convincing evidence that infants come into the world with techniques for segmenting speech into parts, and it is these constraints that allow them later to learn the meaning-bearing units of language. Many models of perception are

shifting away from generic, general-purpose neural networks and toward highly structured, constrained networks that have greater learning and generalization potential because of their preexisting organization (Regier 1996). Early constraints on perception serve to bootstrap the development of more sophisticated percepts. For example, infants seem to be constrained to treat parts that move together as coming from the same object, but this constraint allows them to learn about the color and shape regularities found within objects (Spelke 1990). The preexisting structures that provide the basis of perceptual learning may be innate, but also may be the result of earlier learning processes (Elman et al 1996). At any given time, what can be learned depends on what has already been learned; the constraints on perceptual change may themselves evolve with experience.

Despite limits on the generalization, speed, and occurrence of perceptual learning, it remains an important source of human flexibility. Human learning is often divided into perceptual, cognitive, and procedural varieties. These divisions are regrettable, causing fruitful links to be neglected. There are deep similarities between perceptual unitization and chunking in memory, and between perceptual differentiation and association-building (Hall 1991). In many cases, perceptual learning involves acquiring new procedures for actively probing one's environment (Gibson 1969), such as learning procedures for efficiently scanning the edges of an object (Hochberg 1997, Salapatek & Kessen 1973). Perhaps the only reason to selectively highlight perceptual learning is to stress that flexible and consistent responses often involve adjusting initial representations of stimuli. Perceptual learning exerts a profound influence on behavior precisely because it occurs early during information processing and thus shifts the foundation for all subsequent processes.

In her 1991 preface to her 1963 *Annual Review of Psychology* article, Gibson laments, "I wound up pointing out the need for a theory and the prediction that 'more specific theories of perceptual learning are on the way.' I was wrong there—the cognitive psychologists have seldom concerned themselves with perceptual learning" (Gibson 1991, p. 322). The reviewed research suggests that this quote is too pessimistic; there has been much progress on theories of the sort predicted by Gibson in 1963. These theories are receiving convergent support from several disciplines. Many of the concrete proposals for implementing mechanisms of perceptual learning come from neural and computer sciences. Traditional disciplinary boundaries will have to be crossed for a complete account, and considering the field in terms of underlying mechanisms of adaptation (e.g. attention weighting, imprinting, differentiation, and unitization) rather than domains (e.g. expertise, psychophysics, development, and cross-cultural comparison) will perhaps result in more unified and principled accounts of perceptual learning.

Note: Further information about the topic and author can be accessed via <http://cognitron.psych.indiana.edu/>

Visit the *Annual Reviews* home page at
<http://www.AnnualReviews.org>.

Literature Cited

- Allard T, Clark SA, Jenkins WM, Merzenich MM. 1991. Reorganization of somatosensory area 3b representation in adult Owl Monkeys after digital syndactyly. *J. Neurophysiol.* 66:1048–58
- Allen SW, Brooks LR. 1991. Specializing the operation of an explicit rule. *J. Exp. Psychol. Gen.* 120:3–19
- Anderson JA, Silverstein JW, Ritz SA, Jones RS. 1977. Distinctive features, categorical perception, and probability learning: some applications of a neural model. *Psychol. Rev.* 84:413–51
- Anderson JR. 1987. Skill acquisition: compilation of weak-method problem solutions. *Psychol. Rev.* 94:192–210
- Aslin RN, Smith LB. 1988. Perceptual development. *Annu. Rev. Psychol.* 39:435–73
- Beale JM, Keil FC. 1995. Categorical effects in the perception of faces. *Cognition* 57: 217–39
- Becker S. 1996. Mutual information maximization: models of cortical self-organization. *Netw. Comput. Neural Syst.* 7:7–31
- Bedford FL. 1993. Perceptual learning. In *The Psychology of Learning and Motivation*, ed. DL Medin, pp. 1–60. San Diego: Academic
- Bedford FL. 1995. Constraints on perceptual learning: objects and dimensions. *Cognition* 54:253–97
- Biederman I, Shiffrar MM. 1987. Sexing day-old chicks: a case study and expert systems analysis of a difficult perceptual-learning task. *J. Exp. Psychol.: Learn. Mem. Cogn.* 13:640–45
- Bingham GP, Schmidt RC, Rosenblum LD. 1995. Dynamics and the orientation of kinematic forms in visual event recognition. *J. Exp. Psychol.: Hum. Percept. Perform.* 21:1473–93
- Brooks LR, Norman GR, Allen SW. 1991. Role of specific similarity in a medical diagnostic task. *J. Exp. Psychol. Gen.* 120: 278–87
- Burns B, Shepp BE. 1988. Dimensional inter- actions and the structure of psychological space: the representation of hue, saturation, and brightness. *Percept. Psychophys.* 43:494–507
- Burns EM, Ward WD. 1978. Categorical perception—phenomenon or epiphenomenon: evidence from experiments in the perception of melodic musical intervals. *J. Acoust. Soc. Am.* 63:456–68
- Calder AJ, Young AW, Perrett DI, Etcoff NL, Rowland D. 1996. Categorical perception of morphed facial expressions. *Vis. Cogn.* 3:81–117
- Czerwinski M, Lightfoot N, Shiffrin RM. 1992. Automatization and training in visual search. *Am. J. Psychol.* 105:271–315
- de Sa VR, Ballard DH. 1997. Perceptual learning from cross-modal feedback. See Goldstone et al 1997, pp. 309–52
- Diamond R, Carey S. 1986. Why faces are and are not special: an effect of expertise. *J. Exp. Psychol. Gen.* 115:107–17
- Doane SM, Alderton DL, Sohn YW, Pellegrino JW. 1996. Acquisition and transfer of skilled performance: Are visual discrimination skills stimulus specific? *J. Exp. Psychol. Gen.* 22:1218–48
- Edelman GM. 1987. *Neural Darwinism: The Theory of Neronal Group Selection*. New York: Basic Books
- Edelman S, Intrator N. 1997. Learning as extraction of low-dimensional representations. See Goldstone et al 1997, pp. 353–80
- Eimas PD. 1994. Categorization in early infancy and the continuity of development. *Cognition* 50:83–93
- Eimas PD. 1997. Infant speech perception: processing characteristics, representational units, and the learning of words. See Goldstone et al 1997, pp. 127–70
- Eimas PD, Siqueland ER, Jusczyk PW, Vigorito J. 1971. Speech perception in infants. *Science* 171:303–6
- Elbert T, Pantev C, Wienbruch C, Rockstroh B, Taub E. 1995. Increased cortical represen-

- tation of the fingers of the left hand in string players. *Science*. 270:305–7
- Elman JL, Bates EA, Johnson MH, Karmiloff-Smith A, Parisi D, Plunkett K. 1996. *Rethinking Innateness*. Cambridge, MA: MIT Press
- Fahle M, Edelman S. 1993. Long-term learning in vernier acuity: effects of stimulus orientation, range and of feedback. *Vis. Res.* 33:397–412
- Fahle M, Edelman S, Poggio T. 1995. Fast perceptual learning in hyperacuity. *Vis. Res.* 35:3003–13
- Fahle M, Morgan M. 1996. No transfer of perceptual learning between similar stimuli in the same retinal position. *Curr. Biol.* 6: 292–97
- Farah MJ. 1990. *Visual Agnosia: Disorders of Object Recognition and What They Tell Us About Normal Vision*. Cambridge, MA: MIT Press
- Farah MJ. 1992. Is an object an object an object? Cognitive and neuropsychological investigations of domain-specificity in visual object recognition. *Curr. Dir. Psychol. Sci.* 1:164–69
- Fox E. 1995. Negative priming from ignored distractors in visual selection: a review. *Psychonom. Bull. Rev.* 2:145–73
- Garraghty PE, Kaas JH. 1992. Dynamic features of sensory and motor maps. *Curr. Opin. Neurobiol.* 2:522–27
- Gauthier I, Tarr MJ. 1997. Becoming a “Greeble” expert: exploring mechanisms for face recognition. *Vis. Res.* 37:1673–82
- Geoghegan WH. 1976. Polytopy in folk biological taxonomies. *Am. Ethnol.* 3:469–80
- Gibson EJ. 1969. *Principles of perceptual learning and development*. New York: Appleton-Century-Crofts
- Gibson EJ. 1991. *An Odyssey in Learning and Perception*. Cambridge, MA: MIT Press
- Gibson EJ, Walk RD. 1956. The effect of prolonged exposure to visually presented patterns on learning to discriminate them. *J. Compar. Physiol. Psychol.* 49:239–42
- Gibson JJ, Gibson EJ. 1955. Perceptual learning: differentiation or enrichment? *Psychol. Rev.* 62:32–41
- Gilbert CD. 1996. Plasticity in visual perception and physiology. *Curr. Opin. Neurobiol.* 6:269–74
- Goldstone RL. 1994. Influences of categorization on perceptual discrimination. *J. Exp. Psychol. Gen.* 123:178–200
- Goldstone RL. 1995. Effects of categorization on color perception. *Psychol. Sci.* 6: 298–304
- Goldstone RL. 1996. Isolated and interrelated concepts. *Mem. Cogn.* 24:608–28
- Goldstone RL, Schyns PG, Medin DL, eds. 1997. *The Psychology of Learning and Motivation*. San Diego: Academic
- Grossberg S. 1984. Unitization, automaticity, temporal order, and word recognition. *Cogn. Brain Theory* 7:263–83
- Grossberg S. 1991. Nonlinear neural networks: principles, mechanisms, and architectures. In *Pattern Recognition by Self-Organizing Neural Networks*, ed. GA Carpenter, S Grossberg, pp. 36–109. Cambridge, MA: MIT Press
- Haider H, Frensch PA. 1996. The role of information reduction in skill acquisition. *Cogn. Psychol.* 30:304–37
- Hall G. 1991. *Perceptual and Associative Learning*. Oxford: Clarendon
- Harnad S, ed. 1987. *Categorical Perception*. Cambridge: Cambridge Univ. Press
- Harnad S, Hanson SJ, Lubin J. 1995. Learned categorical perception in neural nets: Implications for symbol grounding. In *Symbolic Processors and Connectionist Network Models in Artificial Intelligence and Cognitive Modelling: Steps Toward Principled Integration*, ed. V Honavar, L Uhr, pp. 191–206. Boston: Academic
- Helson H. 1948. Adaptation-level as a basis for a quantitative theory of frames of reference. *Psychol. Rev.* 55:297–313
- Herrnstein RJ. 1990. Levels of stimulus control: a functional approach. *Cognition* 37: 133–66
- Hillyard HC, Kutas M. 1983. Electrophysiology of cognitive processes. *Annu. Rev. Psychol.* 34:33–61
- Hinton G, Williams K, Revow M. 1992. Adaptive elastic models for handprinted character recognition. In *Advances in Neural Information Processing Systems, IV*, ed. J Moody, S Hanson, R Lippmann, pp. 341–76. San Mateo, CA: Morgan Kaufmann
- Hochberg J. 1997. The affordances of perceptual inquiry: pictures are learned from the world, and what that fact might mean about perception quite generally. See Goldstone et al 1997, pp. 15–44
- Hock HS, Webb E, Cavedo LC. 1987. Perceptual learning in visual category acquisition. *Mem. Cogn.* 15:544–56
- Hofstadter D. 1995. *Fluid Concepts and Creative Analogies*. New York: Basic Books
- Honey RC, Hall G. 1989. Acquired equivalence and distinctiveness of cues. *J. Exp. Psychol. Anim. Behav. Proc.* 15:338–46
- Howells TH. 1944. The experimental development of color-tone synesthesia. *J. Exp. Psychol.* 34:87–103
- Hughes B, Epstein W, Schneider S, Dudock A.

1990. An asymmetry in transmodal perceptual learning. *Percept. Psychophys.* 48: 143–50
- Hunn ES. 1982. The utilitarian factor in folk biological classification. *Am. Anthropol.* 84:830–47
- Intrator N. 1994. Feature extraction using an unsupervised neural network. *Neural Comput.* 4:98–107
- Kaas JH. 1991. Plasticity of sensory and motor maps in adult mammals. *Annu. Rev. Neurosci.* 14:137–67
- Karni A, Sagi D. 1991. Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proc. Natl. Acad. Sci. USA* 88:4966–70
- Karni A, Sagi D. 1993. The time course of learning a visual skill. *Nature* 365:250–52
- Kemler DG. 1983. Holistic and analytic modes in perceptual and cognitive development. In *Perception, Cognition, and Development: Interactional Analyses*, ed. TJ Tighe, BE Shepp, pp. 77–101. Hillsdale, NJ: Erlbaum
- Kolb B. 1995. *Brain Plasticity and Behavior*. NJ: LEA
- Kohonen T. 1995. *Self-Organizing Maps*. Berlin: Springer-Verlag
- Kolers PA, Roediger HL. 1984. Procedures of mind. *J. Verbal Learn. Verbal Behav.* 23: 425–49
- Kuhl PK, Miller JD. 1978. Speech perception by the chinchilla: identification functions for synthetic VOT stimuli. *J. Acoust. Soc. Am.* 63:905–17
- LaBerge D. 1973. Attention and the measurement of perceptual learning. *Mem. Cogn.* 1:268–76
- Lawrence DH. 1949. Acquired distinctiveness of cue. I. Transfer between discriminations on the basis of familiarity with the stimulus. *J. Exp. Psychol.* 39:770–84
- Levin DT. 1996. Classifying faces by race: the structure of face categories. *J. Exp. Psychol.: Learn. Mem. Cogn.* 22:1364–82
- Levinson SC. 1996. Relativity in spatial conception and description. In *Rethinking Linguistic Relativity*, ed. J Gumperz, SC Levinson, pp. 177–202. Cambridge: Cambridge Univ. Press
- Liberman AM, Harris KS, Eimas PD, Lisker L, Bastian J. 1957. The discrimination of speech sounds within and across phoneme boundaries. *J. Exp. Psychol.* 61: 379–88
- Lively SE, Logan JS, Pisoni DB. 1993. Training Japanese listeners to identify English /r/ and /l/. II. The role of phonetic environment and talker variability in learning new perceptual categories. *J. Acoust. Soc. Am.* 94:1242–55
- Livingston KR, Andrews JK. 1995. On the interaction of prior knowledge and stimulus structure in category learning. *Q. J. Exp. Psychol. Hum. Exp. Psychol.* 48A:208–36
- Logan GD. 1988. Toward an instance theory of automatization. *Psychol. Rev.* 95: 492–527
- Logan GD. 1994. Spatial attention and the apprehension of spatial relations. *J. Exp. Psychol.: Hum. Percept. Perform.* 20: 1015–36
- Logan GD, Taylor SE, Etherton JL. 1996. Attention in the acquisition and expression of automaticity. *J. Exp. Psychol.: Learn. Mem. Cogn.* 22:620–38
- Logothetis NK, Pauls J, Poggio T. 1995. Shape representation in the inferior temporal cortex of monkeys. *Curr. Biol.* 5: 552–63
- Lubow RE, Kaplan O. 1997. Visual search as a function of type of prior experience with target and distractor. *J. Exp. Psychol.: Hum. Percept. Perform.* 23:14–24
- Luce RD, Green DM, Weber DL. 1976. Attention bands in absolute identification. *Percept. Psychophys.* 20:49–54
- Mackintosh NJ. 1974. *The Psychology of Animal Learning*. London: Academic
- Malt BC. 1995. Category coherence in cross-cultural perspective. *Cogn. Psychol.* 29: 85–148
- Mandler JM, Bauer PJ, McDonough L. 1991. Separating the sheep from the goats: differentiating global categories. *Cogn. Psychol.* 23:263–98
- McGaugh JL, Bermudez-Rattoni F, Prado-Alcala RA. 1995. *Plasticity in the Central Nervous System*. Mahwah, NJ: LEA
- Melara RD, Marks LE. 1990. Dimensional interactions in language processing: investigating directions and levels of crosstalk. *J. Exp. Psychol.: Learn. Mem. Cogn.* 16: 539–54
- Melcher JM, Schooler JW. 1996. The misremembrance of wines past: verbal and perceptual expertise differentially mediate verbal overshadowing of taste memory. *J. Mem. Lang.* 35:231–45
- Miikkulainen R, Bednar JA, Choe Y, Sirosh J. 1997. Self-organization, plasticity, and low-level visual phenomena in a laterally connected map model of primary visual cortex. See Goldstone et al 1997, pp. 257–308
- Moscovici S, Personnaz B. 1991. Studies in social influence. VI. Is Lenin orange or red? Imagery and social influence. *Eur. J. Soc. Psychol.* 21:101–18
- Mozer MC, Zemel RS, Behrmann M, Williams CKI. 1992. Learning to segment im-

- ages using dynamic feature binding. *Neural Comput.* 4:650–66
- Myles-Worsley M, Johnston WA, Simons MA. 1988. The influence of expertise on X-ray image processing. *J. Exp. Psychol.: Learn. Mem. Cogn.* 14:553–57
- Nosofsky RM. 1986. Attention, similarity, and the identification-categorization relationship. *J. Exp. Psychol. Gen.* 115:39–57
- Nosofsky RM. 1991. Tests of an exemplar model for relating perceptual classification and recognition memory. *J. Exp. Psychol.: Hum. Percept. Perform.* 17:3–27
- Obermayer K, Sejnowski T, Blasdel GG. 1995. Neural pattern formation via a competitive Hebbian mechanism. *Behav. Brain Res.* 66:161–67
- O'Hara W. 1980. Evidence in support of word unitization. *Percept. Psychophys.* 27:390–402
- O'Toole AJ, Peterson J, Deffenbacher KA. 1996. An 'other-race effect' for categorizing faces by sex. *Perception* 25:669–76
- Palmeri TJ. 1997. Exemplar similarity and the development of automaticity. *J. Exp. Psychol.: Learn. Mem. Cogn.* 23:324–54
- Palmeri TJ, Goldinger SD, Pisoni DB. 1993. Episodic encoding of voice attributes and recognition memory for spoken words. *J. Exp. Psychol.: Learn. Mem. Cogn.* 19:309–28
- Pearce JM. 1987. A model for stimulus generalization in Pavlovian conditioning. *Psychol. Rev.* 94:61–73
- Peron RM, Allen GL. 1988. Attempts to train novices for beer flavor discrimination: a matter of taste. *J. Gen. Psychol.* 115:403–18
- Perrett DI, Smith PAJ, Potter DD, Mistlin AJ, Head AD, Jeeves MA. 1984. Neurones responsive to faces in the temporal cortex: studies of functional organization, sensitivity to identity and relation to perception. *Hum. Neurobiol.* 3:197–208
- Peterson MA, Gibson BS. 1994. Must figure-ground organization precede object recognition? An assumption in peril. *Psychol. Sci.* 5:253–59
- Pick HL. 1992. Eleanor J. Gibson: learning to perceive and perceiving to learn. *Dev. Psychol.* 28:787–94
- Pisoni DB, Aslin RN, Perey AJ, Hennessy BL. 1982. Some effects of laboratory training on identification and discrimination of voicing contrasts in stop consonants. *J. Exp. Psychol.: Hum. Percept. Perform.* 8:297–314
- Poggio T, Edelman S. 1990. A network that learns to recognize three-dimensional objects. *Nature* 343:263–66
- Poggio T, Fahle M, Edelman S. 1992. Fast perceptual learning in visual hyperacuity. *Science* 256:1018–21
- Posner MI, Keele SW. 1968. On the genesis of abstract ideas. *J. Exp. Psychol.* 77:353–63
- Quittner AL, Smith LB, Osberger MJ, Mitchell TV, Katz DB. 1994. The impact of audition on the development of visual attention. *Psychol. Sci.* 5:347–53
- Recanzone GH, Merzenich MM, Jenkins WM. 1992. Frequency discrimination training engaging a restricted skin surface results in an emergence of a cutaneous response zone in cortical area 3a. *J. Neurophysiol.* 67:1057–70
- Recanzone GH, Schreiner CE, Merzenich MM. 1993. Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *J. Neurosci.* 13:87–103
- Reed E. 1996. *Encountering the World: Toward an Ecological Psychology*. New York: Oxford Univ. Press
- Regier T. 1996. *The Human Semantic Potential*. Cambridge, MA: MIT Press
- Repp BH, Liberman AM. 1987. Phonetic category boundaries are flexible. See Harnad 1987, pp. 89–112
- Rock I. 1985. Perception and knowledge. *Acta Psychol.* 59:3–22
- Rudy JW, Keither JR, Georgen K. 1993. The effect of age on children's learning of problems that require a configural association solution. *Dev. Psychobiol.* 26:171–84
- Rumelhart DE, Hinton GE, Williams RJ. 1986. Learning internal representations by back-propagating errors. *Nature* 323:533–36
- Rumelhart DE, Zipser D. 1985. Feature discovery by competitive learning. *Cogn. Sci.* 9:75–112
- Saarinen J, Levi DM. 1995. Perceptual learning in vernier acuity: What is learned? *Vis. Res.* 35:519–27
- Sagi D, Tanne D. 1994. Perceptual learning: learning to see. *Curr. Opin. Neurobiol.* 4:195–99
- Saiki J, Hummel JE. 1996. Attribute conjunctions and the part configuration advantage in object category learning. *J. Exp. Psychol.: Learn. Mem. Cogn.* 22:1002–19
- Salapatek P, Kessen W. 1973. Prolonged investigation of a plane geometric triangle by the human newborn. *J. Exp. Child Psychol.* 15:22–29
- Salasoo A, Shiffrin RM, Feustel TC. 1985. Building permanent memory codes: codification and repetition effects in word

- identification. *J. Exp. Psychol. Gen.* 114: 50–77
- Samuel AG. 1981. Phonemic restoration: insights from a new methodology. *J. Exp. Psychol. Gen.* 110:474–94
- Schacter DL. 1987. Implicit memory: history and current status. *J. Exp. Psychol.: Learn. Mem. Cogn.* 13:501–18
- Schmidhuber J, Eldracher M, Foltin B. 1996. Semilinear predictability minimization produces well-known feature detectors. *Neural Comput.* 8:773–86
- Schyns PG, Goldstone RL, Thibaut J. 1998. Development of features in object concepts. *Behav. Brain Sci.* In press
- Schyns PG, Murphy GL. 1994. The ontogeny of part representation in object concepts. In *The Psychology of Learning and Motivation*, ed. DL Medin, 31:305–54. San Diego: Academic
- Schyns PG, Rodet L. 1997. Categorization creates functional features. *J. Exp. Psychol.: Learn. Mem. Cogn.* 23:681–96
- Sekuler AB, Palmer SE, Flynn C. 1994. Local and global processes in visual completion. *Psychol. Sci.* 5:260–67
- Shapiro PN, Penrod SD. 1986. Meta-analysis of face identification studies. *Psychol. Bull.* 100:139–56
- Shiffrin RM, Lightfoot N. 1997. Perceptual learning of alphanumeric-like characters. See Goldstone et al 1997, pp. 45–82
- Shiffrin RM, Schneider W. 1977. Controlled and automatic human information processing. II. Perceptual Learning, automatic attending and a general theory. *Psychol. Rev.* 84:127–90
- Shiu L, Pashler H. 1992. Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Percept. Psychophys.* 52:582–88
- Sinha P, Poggio T. 1996. Role of learning in three-dimensional form perception. *Nature* 384:460–63
- Sireteanu R, Rettenbach R. 1995. Perceptual learning in visual search: fast, enduring, but nonspecific. *Vis. Res.* 35:2037–43
- Smith EE, Haviland SE. 1972. Why words are perceived more accurately than nonwords: inference versus unitization. *J. Exp. Psychol.* 92:59–64
- Smith LB. 1989a. From global similarity to kinds of similarity: the construction of dimensions in development. In *Similarity and Analogical Reasoning*, ed. S Vosniadu, A Ortony, pp. 146–78. Cambridge: Cambridge Univ. Press
- Smith LB. 1989b. A model of perceptual classification in children and adults. *Psychol. Rev.* 96:125–44
- Smith LB, Evans P. 1989. Similarity, identity, and dimensions: perceptual classification in children and adults. In *Object Perception: Structure and Process*, ed. BE Shepp, S Ballesteros, pp. 325–56. Hillsdale, NJ: Erlbaum
- Smith LB, Gasser M, Sandhofer C. 1997. Learning to talk about the properties of objects: a network model of the development of dimensions. See Goldstone et al 1997, pp. 219–56
- Smith LB, Sera M. 1992. A developmental analysis of the polar structure of dimensions. *Cogn. Psychol.* 24:99–142
- Spelke ES. 1990. Principles of object perception. *Cogn. Sci.* 14:29–56
- Stein BE, Wallace MT. 1996. Comparisons of cross-modality integration in midbrain and cortex. *Prog. Brain Res.* 112:289–99
- Stiles J, Tada WL. 1996. Developmental change in children's analysis of spatial patterns. *Dev. Psychol.* 32:951–70
- Strange W, Jenkins JJ. 1978. Role of linguistic experience in the perception of speech. In *Perception and Experience*, ed. RD Walk, HL Pick Jr, pp. 125–69. New York: Plenum
- Tanaka J, Gauthier I. 1997. Expertise in object and face recognition. See Goldstone et al 1997, pp. 83–126
- Tanaka J, Taylor M. 1991. Object categories and expertise: Is the basic level in the eye of the beholder? *Cogn. Psychol.* 23:457–82
- Tarr MJ. 1995. Rotating objects to recognize them: a case study on the role of viewpoint dependency in the recognition of three-dimensional objects. *Psychon. Bull. Rev.* 2:55–82
- Tipper SP. 1992. Selection for action: the role of inhibitory mechanisms. *Curr. Dir. Psychol. Sci.* 1:105–9
- Treisman AM, Gelade G. 1980. A feature-integration theory of attention. *Cogn. Psychol.* 12:97–136
- Ullman S. 1989. Aligning pictorial descriptions: an approach to object recognition. *Cognition* 32:193–254
- Valentine T. 1991. A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Q. J. Exp. Psychol.: Hum. Exp. Psychol.* 43:161–204
- Wang Q, Cavanagh P, Green M. 1994. Familiarity and pop-out in visual search. *Percept. Psychophys.* 56:495–500
- Weinberger NM. 1993. Learning-induced changes of auditory receptive fields. *Curr. Opin. Neurobiol.* 3:570–77
- Werker JF, Lalonde CE. 1988. Cross-language speech perception: initial capabilities and

- developmental change. *Dev. Psychol.* 24: 672–83
- Werker JF, Tees RC. 1984. Cross-language speech perception: evidence for perceptual reorganization during the first year of life. *Infant Behav. Dev.* 7:49–63
- Wisniewski EJ, Medin DL. 1994. On the interaction of theory and data in concept learning. *Cogn. Sci.* 18:221–81
- Wytenbach RA, May ML, Hoy RR. 1996. Categorical perception of sound frequency by crickets. *Science* 273:1542–44
- Zatorre RJ, Halpern AR. 1979. Identification, discrimination, and selective adaptation of simultaneous musical intervals. *Percept. Psychophys.* 26:384–95
- Zohary E, Celebrini S, Britten KH, Newsome WT. 1994. Plasticity that underlies improvement in perceptual performance. *Science* 263:1289–92