Perceptual Learning and Human Expertise

Philip J. Kellman[1,3] and **Patrick Garrigan**[2]

[1]Department of Psychology, University of California, Los Angeles, Los Angeles, CA 90095-1563

[2]Department of Psychology, Saint Joseph's University, Philadelphia, PA 19131

Abstract

We consider perceptual learning -- experience-induced changes in the way perceivers extract information. Often neglected in scientific accounts of learning and in instruction, perceptual learning is a fundamental contributor to human expertise and is likely a crucial contributor in domains where humans show remarkable levels of attainment, such as chess, music, and mathematics. In Section II, we give a brief history and discuss the relation of perceptual learning to other forms of learning. We consider in Section III several specific phenomena, illustrating the scope and characteristics of perceptual learning, including both *discovery* and *fluency* effects. We describe abstract perceptual learning, in which structural relationships are discovered and recognized in novel instances that do not share constituent elements or basic features. In Section IV, we consider primary concepts that have been used to explain and model perceptual learning, including receptive field change, selection, and relational recoding. In Section V, we consider the scope of perceptual learning, contrasting recent research, focused on simple sensory discriminations, with earlier work that emphasized extraction of invariance from varied instances in more complex tasks. Contrary to some recent views, we argue that perceptual learning should not be confined to changes in early sensory analyzers. Phenomena at various levels, we suggest, can be unified by models that emphasize discovery and selection of relevant information. In a final section, we consider the potential role of perceptual learning in educational settings. Most instruction emphasizes facts and procedures that can be verbalized, whereas expertise depends heavily on implicit pattern recognition and selective extraction skills acquired through perceptual learning. We consider reasons why perceptual learning has not been systematically addressed in traditional instruction, and we describe recent successful efforts to create a technology of perceptual learning in areas such as aviation, mathematics, and medicine. Research in perceptual learning promises to advance scientific accounts of learning, and perceptual learning technology may offer similar promise in improving education.

Keywords

[3]Corresponding author: kellman @cognet.ucla.edu. Correspondence concerning this article should be addressed to Philip J. Kellman, Department of Psychology, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, CA, 90095-1563. Email correspondence may be sent to < kellman@cognet.ucla.edu>.

I. Introduction

On a good day, the best human chess grandmaster can beat the world's best chess-playing computer. The computer program is no slouch; every second, it examines upwards of 200 million possible moves. Its makers incorporate sophisticated methods for evaluating positions, and they implement strategies based on advice from grandmaster consultants. Yet, not even this formidable array of techniques gives the computer a clear advantage over the best human player.

If chess performance were based on raw search, the human would not present the slightest problem for the computer. Estimates of human search through possible moves in chess suggest that even the best players examine on the order of 4 possible move sequences, each about 4 plies deep (where a ply is a pair of turns by the two sides). That estimate is per *turn*, not per second, and a single turn may take many seconds. Assuming the computer were limited to 10 sec of search per turn, the human would be at a disadvantage of about 1,999,999,984 moves searched per turn.

Given this disparity, how is it possible for the human to outplay the machine? The accomplishment suggests information processing abilities of remarkable power but mysterious nature. Whatever the human is doing, it is, at its best, roughly equivalent to 2 billion moves per sec of raw search. "Magical" would not seem too strong a description for such abilities.

We have not yet said what abilities these are, but before doing so, we add one more striking observation. Biological systems often display remarkable structures and capacities that have emerged as evolutionary adaptations to serve particular functions. Compared to machines that fly, for example, the capabilities of a dragonfly or hummingbird (or even the lowly mosquito) are astonishing. Yet the information processing capabilities we are considering may be seen as all the more remarkable because they do not appear to be adaptations specialized for one particular task. We did not evolve to play chess. In other words, it is likely that human attainments in chess are consequences of highly general abilities that contribute to learned expertise in many domains. Such abilities may have evolved for more ecological tasks, but they are of such power and generality that humans can become remarkably good in almost any domain involving complex structure.

What abilities are these? They are abilities of *perceptual learning*. The effects we are describing arise from experience-induced changes in the way perceivers pick up information. With practice in any domain, humans become attuned to the relevant features and structural relations that define important classifications, and over time we come to extract these with increasing selectivity and fluency. As a contrast, consider: Most artificial sensing devices that exist, or those we might envision, would have fixed characteristics. If they functioned properly, their performance on the 1000th trial of picking up some information would closely resemble their performance on the first trial. Not so in human perception. Rather, our extraction of information changes adaptively to optimize particular tasks. A large and growing research literature suggests that such changes are pervasive in perception and that they profoundly affect tasks from the pickup of minute sensory detail to

the extraction of complex and abstract relations that underwrite symbolic thought. Perceptual learning thus furnishes a crucial basis of human expertise, from accomplishments as commonplace as skilled reading to those as rarified as expert air traffic control, radiological diagnosis, grandmaster chess, and creative scientific insight.

In this paper, we give an overview of perceptual learning, a long-neglected area of learning, both in scientific theory and research, as well as in educational practice. Our consideration of perceptual learning will proceed as follows. In the second section, we provide some brief historical background on perceptual learning and some taxonomic considerations, contrasting and relating it to other types of learning. In the third section, we consider some instructive examples of perceptual learning, indicating its influence in a range of levels and tasks, and arguing that the information processing changes brought about by perceptual learning can be usefully categorized as *discovery* and *fluency* effects. In the fourth section, we consider explanations and modeling concepts for perceptual learning, and we use this information to consider the scope of perceptual learning in the fifth section. As its role and scope in human expertise become clearer, its absence from conventional instructional settings becomes more paradoxical. In a final section, we discuss these issues and the potential for improving education by using perceptual learning techniques.

II. Perceptual Learning in Context

Perceptual Learning and Taxonomies of Learning

Perceptual learning can be defined as "an increase in the ability to extract information from the environment, as a result of experience and practice with stimulation coming from it." (Gibson, 1969, p. 3). With sporadic exceptions, this kind of learning has been neglected in scientific research on learning. Researchers in animal learning have focused on conditioning or associative learning phenomena – connections between responses and stimuli. Most work on human learning and memory has focused on encoding of items in memory (declarative knowledge) or learning sequences of actions (procedural learning).

Perceptual learning is not encompassed by any of these categories. It works synergistically with them all, so much so that it often comprises a missing link, concealed in murky background issues of other learning research. In stimulus-response approaches to animal and human learning, it is axiomatic that the "stimulus" is part of the physical world and can be described without reference to internal variables in the organism. Used in this way "stimulus" omits a set of thorny issues. For an organism, a physical event is not a stimulus unless it is detected. And what kind of stimulus it is will depend on which properties are registered and how it is classified. Like the sound made (or not) by the proverbial tree falling in the forest, "stimulus" has two meanings, and they are not interchangeable. The tone or light programmed by the experimenter is a physical stimulus, but whether a psychological stimulus is present and what its characteristics are depends on the organism's perceptual capacities.

When stimuli are chosen to be obvious, work in associative learning can occur without probing the fine points of perception and attention. The problem of perceptual learning, however, is that with experience, the organism's pick-up of information *changes.* In its most

fascinating instances, perceptual learning leads the perceiver to detect and distinguish features, differences, or relations not previously registered at all. Two initially indistinguishable stimuli can come to be readily distinguished, even in basic sensory acuities, such as those tested by your optometrist. In higher level tasks, the novice chess player may be blind to the impending checkmate that jumps out at the expert, and the novice art critic may lack the expert's ability to detect the difference between the brush strokes in a genuine Renoir painting and those in a forgery. Perceptual learning is not the attachment of a stimulus to a response, but rather the discovery of new structure in stimulation.

Perceptual learning is also not procedural knowledge. Some learned visual scanning routines for specialized tasks may be procedural, but much perceptual learning can be shown in improved pickup of information in presentations so brief that no set of fixations or scan pattern could drive the relevant improvements. In general, the relation between perceptual learning and procedures is flexible. The perceptual expertise of an instrument flight instructor, for example, may allow her to notice at a glance that the aircraft has drifted 200 feet above the assigned altitude. If she is flying the plane, the correct procedure would be to lower the nose and descend. If she is instructing a student at the controls, the proper procedure may be a gentle reminder to check the altimeter. In examples like this one, the perceptual information so expertly extracted can be mapped onto various responses. Another difference is that in many descriptions (e.g., Anderson, Corbett, Fincham, Hoffman, & Pelletier, 1992), procedures consist of sets of steps that are conscious, at least initially in learning. Availability to consciousness is often not an obvious characteristic of changes in sensitivity that arise in perceptual learning.

Finally, it should be obvious that perceptual learning does not consist of learning declarative information – facts and concepts that can be verbalized. Besides the fact that structures extracted by experts in a domain often cannot be verbally explained, the effect of learning is to change the capacity to extract. This idea has been discussed in instructional contexts (Bereiter & Scardamalia, 1998). Most formal learning contexts implicitly follow a "mind as container" metaphor (see section VI below), with learning as the transfer of declarative knowledge – facts and concepts that can be verbalized. Perceptual learning effects involve "mind as pattern recognizer" (Bereiter & Scardamalia, 1998).

Looking back at our example of chess, some readers may be puzzled by our emphasis on perceptual processes in what appears to be a high-level domain involving explicit reasoning and perhaps language. Although reasoning is certainly involved in chess expertise, it is precisely the difficulty of accounting for human competence within the computational limitations of human reasoning that makes chess, and many other domains of expertise, so fascinating. The most straightforward ideas about explicit reasoning in chess are those that have been successfully formalized and implemented in classic artificial intelligence work. Given a position description (a node in a problem space), there are certain allowable moves, and these may be considered in terms of their value via some evaluation function. Search through the space for the best move is computationally unwieldy, but may be aided by heuristics that prune the search tree. But as we noted earlier, human search of this sort is severely limited and easily dwarfed by computer chess programs. Perhaps more fertile efforts to connect these concepts to human chess playing lie in reasoning about positions as

a basis for heuristically guiding search. It is unlikely, however, that the synergy of pattern recognition and reasoning can be explained by explicit symbolic processes, such as those mediated by language. If chess expertise were based on explicit knowledge, the grandmaster consultants to the developers of chess-playing programs would have long since incorporated the relevant patterns used by the best human players. This has not been feasible because much of the relevant pattern knowledge is not verbally accessible. These are some of the reasons that classic studies of chess expertise (e.g., DeGroot, 1965; Chase & Simon, 1973) have pointed to the crucial importance of perception of structure and a more limited role for explicit reasoning (at least relative to our preconceptions!). We believe that similar conclusions apply to many high-level domains of human competence.

The counterintuitive aspects of perception vs. reasoning in expertise derive both from excessive expectations of reasoning and misunderstandings of the nature of perception. These issues are not new. Max Wertheimer, in his classic work Productive Thinking, discussed formal logical and associative approaches to reasoning and argued that neither encompasses what is perhaps the most crucial process: the apprehension of relations. It was the belief of the Gestalt psychologists, such as Wertheimer, that this apprehension is rooted in perception. Although the point is still not sufficiently appreciated, perception itself is abstract and relational. If we see two trees next to each other, one twice as tall as the other, there are many accurate descriptions that may be extracted via perception. The shade of green of the leaves and the texture of the bark are concrete features. But more abstract structure, such as the ratio of the height of one tree to the other, are just as "real," as are informational variables in stimulation that make them perceivable. This is a deeply important point, one that has proven decisive in modern theories of perception (Gibson, 1969, 1979; Marr, 1982; Michotte, 1952). In recovering the connectivity of a non-rigid, moving entity from otherwise meaningless points of light (Johannson, 1973) or perceiving causal relations from stimulus relations (Michotte, 1952), it is obvious that apprehending the abstract structure of objects, arrangements and events is an important, perhaps the most important, goal of perceptual processing.

Still the foregoing analyses leave some questions unanswered. One might ask why humans play chess (and understand chemistry, etc.) whereas animals do not. Some answers to this question are peripheral to our interests here, such as the fact that instructing new players about rules, procedures, and basic strategy is greatly facilitated by language. Here again, however, the fact that one can produce in a short time novice players who can recite the rules and moves of chess flawlessly contrasts with the fact that no verbal instruction suffices to produce grandmasters. More pertinent, however, is the issue of the pickup of abstract structural relations in perception. Might this faculty differ between humans and most animals? Might humans be more disposed to find abstract structure than animal perceivers? Might language and symbolic functions facilitate the salience of information and help guide pattern extraction? These are fascinating possibilities, which we take up below.

Perceptual Learning and the Origins of Perception

The phrase "perceptual learning" has been used in a number of ways. In classical empiricist views of perceptual development, all meaningful perception (e.g., perception of objects,

motion, and spatial arrangement) was held to arise from initially meaningless sensations. Meaningful perception was thought to derive from associations among sensations (e.g., Locke, 1690/1971; Berkeley, 1709/1910; Titchener, 1902) and with action (Piaget, 1952). In this view, most of perceptual development early in life consists of perceptual learning. This point of view, dominant through most of the history of philosophy and experimental psychology, was based primarily on logical arguments (about the ambiguity of visual stimuli) and on the apparent helplessness of human infants in the first 6 months of life. Young infants' lack of coordinated activity was at once an ingredient of a dogma about early perception and at the same time an obstacle to its direct study.

Over the past several decades, experimental psychologists developed methods of testing infant perception directly. Although infants do not *do* much, they perceive quite a lot. And they do deploy visual attention, via eye and head movements. With appropriate techniques, these tendencies, as well as electrophysiological and other methods, can be used to reveal a great deal about early perception (for a review see Kellman & Arterberry, 1998).

What this research has shown is that the traditional empiricist picture of perceptual development is incorrect. Although perception becomes more precise with age and experience, basic capacities of all sorts – such as the abilities to perceive objects, faces, motion, three-dimensional space, the directions of sounds, coordinate the senses in perceiving events, and other abilities -- arise primarily from innate or early-maturing mechanisms (Bushnell, Sai & Mullin, 1989; Held, 1985; Kellman & Spelke, 1983; Meltzoff & Moore, 1977; Slater, Mattock & Brown, 1990).

Despite its lack of viability as an account of early perceptual development, the idea that learning may allow us to attach meaning to initially meaningless sensations has been suggested to characterize perceptual learning throughout the lifespan. In a classic article, Gibson & Gibson (1955) criticized this view and contrasted it with another.

They called the traditional view – that of adding meaning to sensations via associations and past experience – an *enrichment* view. We must enrich and interpret current sensations by adding associated sensations accumulated from prior experience. In Piaget's (1952; 1954) more action-oriented version of this account, it is also the association of perception and action that leads to meaningful perception of objects, space and events.

The Gibsons noticed a curious fact about enrichment views: The more learning occurs, the less perception will be in correspondence with the actual information present in a situation. This is the case because with more enrichment, the current stimuli play a smaller role in determining the percept. More enrichment means more reliance on previously acquired information. Such a view has been held by generations of scholars who have characterized perception as a construction (Berkeley, 1709/1910; Locke, 1690/1971), an hypothesis (Gregory, 1972), an inference (Brunswik, 1956) or an act of imagination based on past experience (Helmholtz, 1864/1962).

The Gibsons suggested that the truth may be quite the opposite. Experience might make perceivers *better* at using currently available information. With practice perception might become more, not less, in correspondence with the given information. They called this view

of perceptual learning *differentiation*. Differentiation, or the discovery of distinguishing features, may play an important role in various creative and cognitive processes. Through discovery, undifferentiated concepts come to have better-defined boundaries (Jung, 1923). Creativity, it has been argued, involves connecting various well differentiated concepts to unconscious representations of knowledge, like instincts and emotion (Perlovsky, 2006). In any situation, there is a wealth of available information to which perceivers might become sensitive. Learning as differentiation is the discovery and selective processing of the information most relevant to a task. This includes filtering relevant from irrelevant information, but also discovery of higher-order invariants that govern some classification. Perceptual learning is conceived of as a process that leads to task-specific improvement in the extraction of what is really there.

In our review, we focus on this notion of perceptual learning, that is, learning of the sort that constitutes an *improvement* in the extraction of information. Most contemporary perceptual learning research, although varying on other dimensions, fits squarely within the differentiation camp and would be difficult to interpret as enrichment. Although contemporary computational vision approaches that emphasize use of Bayesian priors in determining perceptual descriptions may be considered updated, more quantitative, versions of traditional enrichment views, the use of enrichment as an account of learning by the individual is rare in perceptual learning research today. This is partly due to the consistent emphasis on perceptual learning phenomena that involve improvements in sensitivity (rather than changes in bias). Also, enrichment is not much reflected in contemporary perceptual learning work because there is little evidence of perception being substantially influenced by accumulation of priors ontogenetically (i.e., during learning by the individual). Current Bayesian approaches to perception often suggest that priors have been acquired during evolutionary time (e.g., Purves, Williams, Nundy, & Lotto, 2004). This emphasis on evolutionary origins is much more consistent with what is known about early perceptual development (Kellman & Arterberry, 1998).

The contemporary focus on differentiation in perceptual learning can also understood in terms of signal detection theory (SDT). SDT is a psychological model of how biological systems can detect signals in the presence of noise. It is closely related to the concept of signal to noise ratio in telecommunications (e.g., Cover & Thomas, 1991), but also considers potential response biases of subjects that can arise due to, for example, the costs associated with failing to detect a signal or mistaking noise for a signal. More specifically, SDT offers mathematical techniques for quantifying sensitivity independent of response bias from perceptual data, where sensitivity represents accuracy in detecting or discriminating and response bias represents the tendencies of an observer to use the available response categories in the presence of uncertainty.

A well-established fact of SDT is that stimulus frequency affects response bias, not sensitivity (e.g., Wickens, 2002). Consider an experiment in which the observer must say red or blue for a stimulus pattern presented on each trial, and the discrimination is difficult (so that accuracy is well below perfect). For an ambiguous stimulus, observers will be more likely to respond blue if "blue" is the correct answer on 90% of trials than if "blue" is the correct answer on 50% of trials. SDT analysis will reveal that sensitivity is the same, and

what has changed is response bias. Such frequency effects contain the essential elements both of associative "enrichment" theories of perceptual learning and of the use of Bayesian priors in generating perceptual responses.

Although there are interesting phenomena associated with such changes in response bias (including optimizing perceptual decisions), perceptual learning research today is much more concerned with the remarkable fact that practice actually changes *sensitivity*. Sensitivity in SDT is a function of the difference between the observer's probability of saying blue given a blue stimulus and the probability of saying "blue" given a red stimulus). The perceptual learning effects we will consider are typically improvements in sensitivity or facility in dealing with information available to the senses. As we will see, such improvements can come in the form of finer discriminations or discovery and selective use of higher-order structure.

A Brief History of Perceptual Learning

This notion of perceptual learning as improvement in the pick-up of information has not been a mainstay in the psychology of learning. Still it has made cameo appearances. William James devoted a section of his Principles of Psychology to "the improvement of discrimination by practice (James, 1890, Vol. I, p. 509)." Clark Hull, the noted mathematical learning theorist, did his dissertation in 1918 on a concept learning experiment in which slightly deformed Chinese characters were used. In each of twelve categories of characters, differing exemplars shared some invariant structural property. Subjects learned to associate a sound (as the name of each category) using 6 instances of each category and were trained to a learning criterion. A test phase with 6 new instances showed that learning had led to the ability to classify novel instances accurately (Hull, 1920). This ability to extract invariance from instances and to respond selectively to classify new instances mark Hull's study as not only a concept formation experiment but as a perceptual learning experiment.

Hull's later contributions to learning theory had little to do with this early work, a bit of an irony, as Eleanor Gibson, who did the most to define the modern field of perceptual learning, was later a student in Hull's laboratory. She found Hull's laboratory preoccupied with conditioning phenomena when she did her dissertation there in 1938 (Szokolszky, 2003). It was Gibson and her students who made perceptual learning a highly visible area of research in the 1950s and 1960s, much of it summarized in a classic review of the field. By the mid-1970s, this area had become quiescent. One reason was that much of the research at the time had shifted focus to infant perception and cognition. Most of this research focused on characterizing basic perceptual capacities of young infants rather than on perceptual learning processes. Research specifically directed at perceptual learning in infancy has become more common recently, however (e.g., Fiser & Aslin, 2002; Saffran, Loman, & Robertson, 2000; Gomez & Gerken, 2000).

In the last decade and a half, there has been an explosion of research on perceptual learning. This newest period of research has distinctive characteristics. It has centered almost exclusively on learning effects involving elementary sensory discriminations. The research has produced remarkable findings of marked improvements on almost any sensory task brought about by discrimination practice. The focus on basic sensory discriminations has

been motivated by interest in cortical plasticity and attempts to connect behavioral and physiological data. This focus differs from earlier work on perceptual learning, as we consider below, and relating perceptual learning work spanning different tasks and levels poses a number of interesting challenges and opportunities. One goal of this review is to highlight phenomena and principles appearing at different levels and in different tasks that point toward a unified understanding of perceptual learning.

III.   Some Perceptual Learning Phenomena

To give some idea of its scope and characteristics, we consider a few examples of perceptual learning phenomena. Our examples come primarily from visual perception. They span a range of levels, from simple sensory discriminations to higher-level perceptual learning effects more relevant to real-world expertise.

Highlighting these examples may also illustrate that perceptual learning phenomena can be organized into two general categories – *discovery* and *fluency* effects (Kellman, 2002). Discovery involves some change in the bases of response: Selecting new information relevant for a task, amplifying relevant information, or suppressing information that is irrelevant are cases of discovery. Fluency effects involve changes, not in what is extracted, but in the ease of extraction. When relevant information has been discovered, continued practice allows its pickup to become less demanding of attention and effort. In the limit, it becomes *automatic* -- relatively insensitive to attentional load (Schneider & Shiffrin, 1977).

The distinction between discovery and fluency is not always clearcut or easily assessed, as these effects typically co-occur in learning. Particular dependent variables may be sensitive to both. Improved sensitivity to a feature or relation would characterize discovery, whereas improved speed, load insensitivity, or dual-task performance would naturally seem to be fluency effects. But improved speed may also be a consequence of discovering a higher-order invariant that provides a shortcut for the task (a discovery effect). Conversely, when performance is measured under time constraints, fluency improvements may influence measures of sensitivity or accuracy, because greater fluency allows more information to be extracted in a fixed time. The distinction between discovery and fluency is nevertheless important; understanding perceptual learning will require accounts of how new bases of response are discovered and how information extraction becomes more rapid and less demanding with practice.

The examples we consider illustrate a variety of perceptual learning effects but do not comprise a comprehensive treatment of phenomena in the field. Also, there are a number of interesting issues in the field that fall outside the scope of this article. These include several factors that affect perceptual learning, including the effects of sleep (e.g. Karni & Sagi, 1993; Stickgold, James, & Hobson, 2000), the importance or irrelevance of feedback (e.g. Shiu & Pashler, 1992; Fahle & Edelman, 1993; Petrov, Dosher, & Lu, 2005), and even whether or not subjects need to be perceptually aware of the stimuli (Watanabe, Nanez, & Sasaki, 2001; Seitz & Watanabe, 2003). For these and other issues not treated here, we refer the interested reader to more inclusive reviews (Gibson, 1969; Goldstone, 1998; Kellman, 2002; Fahle & Poggio, 2002).

Basic Visual Discriminations

What if someone were to tell you that your visual system's maximal ability to resolve fine detail could improve with practice – by orders of magnitude? This may surprise you, as you might expect that basic limits of resolution, such as discriminating two lines with slightly different orientations, would rest on mechanisms that are fixed. Because basic characteristics of resolution affect many higher-level processes, we might expect on evolutionary grounds that observers' efficiency for extracting detail is optimal within biophysical constraints. In fact, this is not the case – not for orientation sensitivity or for a variety of other basic, visual sensitivities. Appropriately structured perceptual learning can produce dramatic changes in basic sensitivities.

Vernier acuity – judging the alignment of two small line segments – is a frequently studied task that reveals the power of perceptual learning. In a typical Vernier task, two lines, one above the other, are shown on each trial, and the observer indicates whether the upper line is to the left or right of the lower line. In untrained subjects, Vernier acuity is already remarkably precise. So precise, in fact, that it is not uncommon for subjects to be able to detect misalignments of less than 10 arc secs (Westheimer & McKee, 1978). Because this detection threshold is smaller than the aperture diameter of one foveal cone photoreceptor, this acuity, along with others, has been labelled a "hyperacuity" (Westheimer, 1975).

Hyperacuity phenomena have posed challenges to vision modelers. Recent modeling efforts suggest that the basic acuity limits of the human visual system are achieved in part by exploiting sources of information that arise due to the non-linear response properties of retinal and cortical cells responding to near-threshold signals in the presence of noise (Hongler, de Meneses, Beyeler, and Jacot, 2003; Hennig, Kerscher, Funke, and Worgotter, 2002). Despite the neural complexity already involved in achieving hyperacuity, with training over thousands of trials, Vernier thresholds *improve* substantially. Saarinen & Levi (1995), for example, found as much as a 6-fold decrease in threshold after 8000 trials of training. Similar improvements have been found for motion discrimination (Ball & Sekuler, 1982) and simple orientation sensitivity (Shiu & Pashler, 1992; Vogels & Orban, 1985).

Visual Search

A task with more obvious ecological relevance is visual search. Finding or discriminating a target object hidden among distractors, or in noise, characterizes a number of tasks encountered by visually guided animals, including detection of food and predators. Studies of perceptual learning in various visual search tasks show that experience leads to robust gains in sensitivity and speed. Karni & Sagi (1993) had observers search for obliquely oriented lines in a field of horizontal lines. The target, always in one quadrant of the visual field, consisted of a set of three oblique lines arrayed vertically or horizontally. (The subject's task was to say on each trial whether the vertical or horizontal configuration was present.) The amount of time needed to reliably discriminate decreased from about 200 ms on session 1 to about 50 ms on session 15. Sessions were spaced 1–3 days apart, and training effects persisted even years later. Data from their study are shown in Figure 1.

The learning observed by Karni & Sagi was specific to location and orientation (of the background elements). This specificity is observed in many perceptual learning tasks using basic discriminations, a small number of unchanging stimulus patterns, and fixed presentation characteristics. The lack of transfer to untrained retinal locations, across eyes, and to different stimulus values is often argued to indicate a low-level locus of learning. Learning that is specific to a retinal position, for example, may involve early cortical levels in which the responses of neural units are retinotopic – mapped to particular retinal positions – rather than units in higher level areas which show some degree of positional invariance. On the other hand, other studies of perceptual learning have shown robust transfer. In visual search studies, Ahissar & Hochstein (2000) showed that learning to detect a single line element hidden in an array of parallel, differently-oriented line segments could generalize to positions at which the target was never presented. Sireteanu & Rettenbach (2000) discussed learning effects in which learning leads serial (sequential) search tasks to become parallel or nearly so; such effects often generalize across eyes, retinal locations, and tasks. Inconsistency of results regarding specificity of transfer have complicated the idea that lack of transfer implies a low-level site of learning. The basic inference that specificity indicates that learning occurs at early processing levels has also been argued to be a fallacy (Mollon & Danilova, 1996). We discuss this issue in relation to general views of perceptual learning in section V.

Automaticity in Search

Earlier work involving perceptual learning in visual search provides a classic example of improvements in fluency. In a series of studies, Schneider & Shiffrin (1977) had subjects judge whether any letters in a target set appeared at the corners of rectangular arrays that appeared in sequence. Attentional load was manipulated by varying the number of items in the target set and the number of items on each card in the series of frames. (Specifically, load was defined as the product of the number of possible target items and the number of possible locations to be searched.) Early in learning, or when targets and distractors were interchangeable across trials, performance was highly load-sensitive: Searching for larger numbers of items or checking more locations took longer. For subjects trained with a consistent mapping between targets and distractors (i.e., an item could appear as one or the other but not both), subjects not only became much more efficient in doing the task, but they came to perform the task equally well over a range of target set sizes and array sizes. The fact that performance became load-insensitive led Schneider & Shiffrin to label this type of performance *automatic* processing.

Unitization in Shape Processing

Some perceptual learning effects on visual search suggest that the representation of targets, distractors, or both, can change due to experience. Consider the results of Sireteanu & Rettenbach (2000), in which a previously serial visual search comes to be performed in parallel as a result of training. The classic interpretation of this result is that training caused the target of the visual search to become represented as a visual primitive, or feature (Treisman & Gelade, 1980). More generally, this result suggests that conjunctions of features can become represented as features themselves, i.e., they become a unit.

Unitization, like "chunking" (Miller 1956), refers to the process of how information can get encoded or processed more efficiently. The idea is that the information capacity of various cognitive functions is fixed and limited. We can, however, learn to utilize the available capacity more efficiently by forming higher-level units based on relations discovered through learning.

Goldstone (2000) studied unitization using simple 2D shapes like those shown in Figure 2. Each shape is formed from 5 separate parts. Subjects were trained to sort these shapes into categories based on the presence of a single component, a conjunction of 5 components, or a conjunction of 5 components in a specific order. Performance, measured by response time, improved with training, and although categorization based on a single feature was fastest, the greatest improvement due to training was for categorization based on 5 parts in a specific order. Goldstone suggests that the improvement was because the shapes came to be represented in a new way. The original representation of the shapes, through training, was replaced with new, more efficient, chunk-like representation formed from parts that were always present in the stimuli, but previously not perceived. A detailed study of response times indicated that the improvement based on unitization was greater than would be predicted based on improvements in processing of individual components. This result furnished strong evidence that new perceptual units were being used.

An interesting issue with this and other sources of evidence for perceptual unitization or chunking is whether basic elements have merely become cemented together perceptually or whether a higher-order invariant has been discovered that spans multiple elements and replaces their separate coding (Gibson, 1969). In the top display in Figure 2, for example, the learner may come to encode relations among positions of peaks or valleys in the contour. Such relations are not defined if individual components are considered separately. It is possible that unitization phenomena in general depend on discovery of higher-order pattern relations, but these ideas deserve further study.

Perceptual Learning in Real-World Tasks

In recent years, perceptual learning research has most often focused on basic sensory discriminations. Both the focus and style of this research contrast with real-world learning tasks, in which it would be rare for a learner to have thousands of trials discriminating two simple displays differing minimally on some elementary attribute. Ecologically, the function of perceptual learning is almost certainly to allow improvements in information extraction from richer, multidimensional stimuli, where even those falling into some category show a range of variation. A requirement of learning in such tasks is that the learner comes to extract invariant properties from instances having variable characteristics. As E. Gibson put it *"It would be hard to overemphasize the importance for perceptual learning of the discovery of invariant properties which are in correspondence with physical variables. The term 'extraction' will be applied to this process, for the property may be buried, as it were, in a welter of impinging stimulation."(Gibson, 1969, p. 81)*.

These aspects of real-world tasks at once indicate why perceptual learning is so important and why it is task-specific. For a given purpose, not all properties of an object are relevant. In fact, for understanding causal relations and structural similarities, finding the *relevant*

structure is a key to thinking and problem solving (Duncker, 1945; Wertheimer, 1957). For different tasks, the relevant properties will vary. For example, in classifying human beings, the relevant properties differ in an employment interview and in an aircraft weight and balance calculation (at least for most jobs). Often, perceptual learning in real-world situations involves discovery of relations and some degree of abstraction. Perceptual learning that ferrets out dimensions and relations crucial to particular tasks underwrites not only remarkable improvements in orientation discrimination but also the expertise that empowers a sommelier or successful day-trader. In understanding human expertise, these factors play a greater role than is usually suspected; conversely, the burden of explanation placed on the learning of facts, rules, or techniques is often exaggerated, and the dependence of these latter processes on perceptual learning goes unnoticed. Some examples may illustrate this idea.

Chess Expertise.—Chess is a fascinating domain in which to study perceptual learning, for several reasons. One is that, as mentioned earlier, humans can reach astonishing levels of expertise. Another is that the differences between middle level players and masters tend not to involve explicit knowledge about chess. DeGroot (1965) and Chase & Simon (1973) studied chess masters, intermediate players, and novices to try to determine how masters perform at such a high level. Masters did not differ from novices or mid-level players in terms of number of moves considered, search strategies, or depth of search. Rather, the differences appeared to be perceptual learning effects – exceptional abilities to encode rapidly and accurately positions and relations on the board. These abilities were tested in experiments in which players were given brief exposures to chess positions and had to recreate them with a set of pieces on an empty board. Relative to intermediate or novice players, a chess master seemed to encode larger chunks, picked up important chess relations in these chunks and required very few exposures to fully recreate a chessboard. One might wonder if masters are simply people with superior visual memory skills – skills that allowed them to excel at chess. A control experiment (Chase & Simon, 1973) suggested that this is not the case. When pieces were placed randomly on a chessboard, a chess master showed no better performance in recreating board positions than a novice or mid-level player. (In fact, there was a slight tendency for the master to perform worse than the others.) These results suggest that the master differs in having advanced skills in extracting structural patterns specific to chess.

The Word Superiority Effect.—It is comforting to know that perceptual expertise is not the sole province of grandmasters, as few of us would participate. Common to almost all adults is a domain in which extensive practice leads to high expertise: fluent reading. An intriguing indicator of the power of perceptual learning effects in reading is the *word superiority effect*. Suppose a participant is asked to judge on each trial whether the letter R or L is presented. Suppose further that the presentations consist of brief exposures and that the time of these exposures is varied. We can use such a method to determine an exposure time that allows the observer to be correct about 80% of the time.

Now suppose that, instead of R or L, the observer's task is to judge whether "DEAR" or "DEAL" is presented on each trial. Again, we find the exposure time that produces 80%

accuracy. If we compare this condition to the single letter condition, we see a remarkable effect first studied by Reicher (1969) and Wheeler (1970). (For an excellent review, see Baron, 1978.) The exposure time allowing 80% performance will be substantially shorter for the words than for the individual letters. Note that the context provided by the word frame provides no information about the correct choice, as the frame is identical for both. In both the single letter and word conditions, the discrimination must be based on the R vs. the L. Somehow, the presence of the irrelevant word context allows the discrimination to be made more rapidly.

One might suppose that the word-superiority effect comes from readers having learned the overall shapes of familiar words. If practice leads to discovery of higher-order structure of the word (beyond the processing of the individual letters), this structure may become more rapidly processed. It turns out that this explanation of the word superiority effect is not correct. The most stunning evidence on the point is that the effect is not restricted to familiar words; it also works for pronounceable nonsense (e.g., GEAL vs. GEAD, but not GKBL vs. GKBD). The fact that most of the effect is preserved with pronounceable nonsense suggests that learning has led to extraction and rapid processing of structural regularities characteristic of English spelling patterns. The idea that the word superiority effect results from perceptual learning of higher-order regularities would predict that beginning readers would not show the word superiority effect. This proves to be the case (Feitelson & Razel, 1984).

Interaction of Fluency and Discovery Processes.—It has long been suspected that fluency and discovery processes interact in the development of expertise. Writing in *Psychological Review* in 1897, Bryan & Harter proposed that automatizing the pick-up of basic structure serves as a foundation for discovering higher-order relationships (Bryan & Harter, 1897). These investigators studied learning in the task of telegraphic receiving. When the measure of characters received per minute (in Morse Code) 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. And for some subjects, after even more practice, a third learning curve ultimately emerged. Each learning curve raised performance to substantially higher levels than before.

What could account for this remarkable topography of learning? When Bryan & 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 & 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 subjects' subjective reports.

The Bryan & Harter results serve as a good example for distinguishing discovery and fluency effects in perceptual learning. Both are involved in the telegraphers' learning. As learning curves are ascending, we may assume that at least some of what is occurring is that learners are discovering new structure that enables better performance. The acquisition of new bases of response (discovery) is clearly shown in the tests Bryan & Harter did with non-meaningful letter and word sequences. What is also noticeable, however, are long periods at asymptotic performance at one level before a new learning curve emerges. In Figure 2 (top curve), such a period continues for nearly a month of practice. Given that performance (words per minute transcribed) is not changing, the information being extracted is unlikely to be changing in such periods. What is changing? Arguably, what is changing is that controlled, attentionally demanding processing is giving way to automatic processing (Schneider & Shiffrin, 1973). If in fact the information being extracted is constant, then this change would appear to be a pure fluency effect. Both the discovery of structure and automatizing its pick-up appear to be necessary to pave the way for discovery of higher level structure.

Although the universality of the phenomenon of three separable learning curves in telegraphic receiving has been questioned (Keller, 1958), Bryan & Harter's data indicate use of higher-order structure by advanced learners and suggest that discovery of such structure is a limited capacity process. Automating the processing of basic structure at one level frees attentional capacity to discover higher level structure, which can in turn become automatic, allowing discovery of even higher level information, and so on. This cycle of discovery and fluency in perceptual learning --- discovering and automating 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. These insights from the classic Bryan & Harter study are reflected in more recent work, such as demonstrations that expertise involves learning across multiple time scales in pure motor skill development (Liu, Mayer-Kress, and Newell, 2006) and automating basic skills in development of higher-level expertise in areas such as reading (Samuels and Flor, 1997) and mathematics (Gagne, 1983). These principles are increasingly reflected in modern application of perceptual learning to training procedures (e.g., Clawson, Healy, Ericcson, and Bourne, 2001). Bryan & Harter's study offers a compelling suggestion about how discovery and fluency processes interact. Their 1897 article ends with a memorable claim: "Automaticity is not genius, but it is the hands and feet of genius."

Abstract Perceptual Learning

Two important tenets of contemporary theories of perception are that perceptual systems are sensitive to higher-order relationships (Gibson, 1979; Marr, 1982), and they produce as outputs abstract descriptions of reality (Marr, 1982; Michotte, 1952), that is, descriptions of physical objects, shapes, spatial relations and events. These modern ideas contrast with pervasive earlier views that sensory systems produce initially meaningless sensations, which acquire meaning only through association and inference.

Sensitivity to abstract relations is reflected as well in perceptual learning. It is probably worth making clear what we mean by abstract, as there are a variety of uses of the term. For our purposes, abstract information may be illustrated as follows. Suppose we hear and encode a melody. Encoding the information that the first note has fundamental frequency f_1 is a concrete encoding. Encoding that the second note has a fundamental frequency that is higher in frequency than the first note by some amount k is encoding a *relation*. Encoding that the second note is higher than the first by an octave is encoding an *abstract relation*. The idea is that the abstract encoding is dependent on a concrete value only conditionally. For an octave, whatever the first frequency is, the second must be twice that. In general, a relation is abstract if it involves binding the value of a variable. (In this example, the first frequency can be any frequency x, but the second must be 2x.) Our notion of abstract information is close to J. Gibson's notion of higher-order invariants (Gibson, 1950).

The importance of relations and abstract relations was crucial in Gestalt Psychology (e.g., Wertheimer, 1923/1938; Koffka, 1935). Our use of a melody borrows from the Gestaltists, as it was a classic example of theirs. If you hear a new melody and remember it, what is it that you have learned? If learning were confined to a concrete, sensory level, the learning would consist of a temporal series of frequencies of sound. But this description would not capture the learning of most listeners. Most of us (except those with "perfect pitch") will not retain the particular frequencies (or more accurately, the sensations of pitch corresponding to those frequencies). Your learning is such that you will seamlessly recognize the melody as the same if you hear it later, even if it has been transposed to a different key. The fact that the melody retains its identity despite changes in all of the constituent frequencies of sound indicates encoding of the melody in terms of abstract relations. Conversely, most hearers would have little ability to identify later the exact key in which the melody was played at first hearing.

In vision, abstract relations are pervasive in our encoding of shape. A miniature plastic replica of an elephant is easily recognizable as an elephant. This ability is more remarkable than it first appears. It involves abstract relations and using these to classify despite novel (and starkly conflicting) concrete features. (No real elephants are made of plastic or fit in the palm of your hand.) Seeing the miniature elephant as appropriately shaped involves abstract relations, such as the proportions of the body, trunk and ears, that are applied to a novel set of concrete size values. Perceptual learning about shape seems to have this property in any situation in which shape invariance is required across changes in constituent elements, size or other changes (what the Gestalt psychologists called "transposition" phenomena).

In the cases of a melody and for some shape perception, abstract relations are salient even in initial encoding (although there is a possibility that such invariance is initially discovered through learning processes). In other contexts, learning processes have a larger challenge in discovering over a longer period some higher-order invariance that determines a classification (Gibson, 1969). Whether the relevant relations "pop out" or are discovered gradually, the pickup of abstract structure is common in human perceptual learning, and it presents challenges in the modeling of perceptual learning.

There has not been much work on perceptual learning of abstract relations. Numerous efforts have involved shape perception (Gibson, 1969), but not much has been done to specifically examine invariance over transformations after learning. Intuitively, work on caricature (Ryan & Schwartz, 1956) or on recognition of people and events in point-light displays (e.g., a person walking has small lights placed on the major joints of the body and is filmed in the dark; Johansson, 1973) are studies of the effects of abstract perceptual learning processes.

Although sensory discrimination tasks used in most contemporary perceptual learning research do not appear to involve abstract perceptual learning, it may be more pervasive than suspected. An interesting example comes from the work of Nagarajan, Blake, Wright, Byl, and Merzenich (1998). They trained subjects on an interval discrimination for vibratory pulses. The sensory (concrete) aspects of this learning would be expected to involve tactile sensations sent to somatosensory cortex. They found that after training, learning transferred to other parts of the trained hand, and to the same position on the contralateral hand. More remarkable, they found that learning also transferred to auditory stimuli with the same intervals. So, what was learned was actually the *time* interval – an abstract relation among inputs, with the particular inputs being incidental. This idea that even apparently low-level tasks may involve relational structure is also reflected in recent work suggesting that perceptual learning may not operate directly on sensory analyzers but only through perceptual-constancy based representations (Garrigan & Kellman, 2008). We discuss that work more thoroughly in section V below.

Research suggests that perceptual learning of abstract relations is a basic characteristic of human perceptual learning from very early in life. Marcus, Vijayan, Bandi Rau & 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, 2000). These findings suggest an early ability to discover abstract relationships, although there is some possibility that speech stimuli are somewhat special in this respect (Saffran & Greipentrog, 2001).

That learning often latches onto abstract relations is important in attaining behaviorally relevant descriptions of our environment: For thought and action, it is often the case that encoding relations and abstract relations is more crucial than encoding sensory particulars (Gibson, 1979; Koffka, 1935; von Hornbostel, 1927). Whether this holds true depends on the task and environment, of course.

IV. Explaining and Modeling Perceptual Learning Phenomena

How have researchers sought to explain and model perceptual learning? These questions are important not only for understanding human performance but for artificial systems as well. Understanding how learners discover invariance among variable instances would have value for creating learning devices as well as explaining human abilities. We currently have no good machine learning algorithms that can learn from several examples to correctly classify new instances of dogs, cats, and squirrels the way 4-year-old humans do routinely. Even

understanding how basic acuities improve with practice presents interesting challenges for modeling. Here we describe several foundational ideas that have been proposed and illustrated by empirical studies or models, or both.

Receptive Field Modification

Aligned with the focus of much research is the idea that low-level perceptual learning might work by modifying the receptive fields of the cells that initially encode the stimulus. For example, individual cells could adapt to become more sensitive to important features, effectively recruiting more cells for a particular purpose, making some cells more specifically tuned for the task at hand, or both. Equivalently, receptive field modification can be thought of as a way to exclude irrelevant information. A detector that is sensitive to two similar orientations might develop narrower sensitivity to facilitate discrimination between the two.

The idea of receptive field modification in early cortical areas would fit with some known properties of perceptual learning. In vision, specificity to retinal location and to particular ranges on stimulus dimensions (e.g., orientation) would be consistent with known properties of cells in early cortical areas (V1 and V2). Consistent with some perceptual learning results, effects of changing a cell's receptive field would be long lasting, compared to other adaptation effects. As a cell's receptive field becomes more specifically tuned, it may also become more resilient to future, experience-induced changes since it would be less broadly sampling the statistics of the environment.

Evidence for receptive field change has been found using single-cell recording techniques in primates. In monkey somatosensory cortex, there is high variability in the size of representations of the digits (Merzenich, et al., 1987). At least some of this variability may be due to effects of experience. Perceptual learning can dramatically alter both the total amount of cortex devoted to a particular mental representation as well as the size of the receptive fields of the individual cells. Recanzone, et al., (1992) trained owl monkeys to discriminate different frequencies of tactile vibration. Following training, the total size of cortex corresponding to the trained digits increased 1.5–3 fold. Other studies have shown that a task requiring fine dexterity, e.g. retrieving food pellets from small receptacles, resulted in a decrease in the size of the receptive fields of cells corresponding to the tips of the digits used in the task (Xerri, et al., 1999). Presumably, smaller receptive fields would lead to finer tactile acuity, and thereby increase the level of performance in retrieving the pellets.

In the auditory domain, Recanzone, Schreiner & Merzenich (1993) showed that monkeys trained on a difficult frequency discrimination improved over several weeks. Mapping of primary auditory cortex after training showed that receptive fields of cortical cells responding to the task-relevant frequencies were larger than before training.

An example of a specific model that attempts to explain learning effects by receptive field change is the vision model of Vernier hyperacuity proposed by Poggio, Fahle, & Edelman (1992). Their model begins with a network of non-oriented units receiving input from simulated photoreceptors. The photoreceptor layer of their model remains static, while the

second layer is an optimized representation of the photoreceptor activity using radial basis functions. This second layer represents cells whose receptive field structure changes to more effectively utilize task-relevant information from the photoreceptors.

Selection and Differentiation

An alternative approach for modeling perceptual learning is the idea of *selection*. Selection describes learning to preferentially use some subset of the information available for making a decision. This notion was at the core of Eleanor Gibson's work, so much so that she often used "differentiation learning" as synonymous with perceptual learning. Summing up the approach, she said:

> It has been argued that what is learned in perceptual learning are distinctive features, invariant relationships, and patterns; that these are available in stimulation; that they must, therefore, be extracted from the total stimulus flux. … From the welter of stimulation constantly impinging on the sensory surfaces of an organism, there must be selection. (Gibson, 1969, p. 119.)

A specific proposal put forth by Gibson was that perceptual learning works by discovering distinguishing (or distinctive) features. Distinguishing features are task-specific: they are those features that provide the contrasts relevant to some classification. Thus, in learning to discriminate objects, the learner will tend to select information along dimensions that distinguish the objects, rather than forming better overall descriptions of the objects (Pick, 1965).

In contemporary perceptual learning work with basic sensory discriminations, the notion of selection is also a viable candidate for explaining many results. Imagine an experiment in which the learner improves at distinguishing two different Gabor patches (or lines) having a slight orientation difference. On each trial the learner must decide which of the two patches was presented. Now envision in cortex an array of orientation-sensitive units responsive at the relevant stimulus positions. Some units are activated by either stimulus; some more by one stimulus than the other; others may be activated by only one stimulus. Suppose the learner's decision initially takes all of these responses into account, giving each equal measure. As learning proceeds, the weights of these inputs are gradually altered. Units that do not discriminate well between the stimuli will be given less weight, whereas those that discriminate strongly will be given more weight. It may even be the case that the best discriminators are not the units that initially give the largest response to either stimulus. For example, suppose the learner is to discriminate lines at 10^o vs. 14^o, and suppose orientation-sensitive units respond somewhat to inputs within $\pm 6^o$ of their preferred orientations. In this simple example, an analyzers at 7^o and 17^o may be more informative than the units centered at 10^o or 14^o, because such analyzers would be activated by only one of the two displays.

Recently, Petrov, Dosher & Lu (2005) presented experimental and modeling results indicating that in a low-level task (orientation discrimination of two Gabor patches in noise), perceptual learning was best described, not in terms of receptive field modification, but by *selective reweighting*. They argued that even in simple discrimination paradigms, perceptual learning might be explained in terms of discovery of which analyzers best perform a

classification and the increased weighting of such analyzers. In their model, encodings at the lowest level remain unchanged.

This approach brings to light an interesting set of issues. Earlier models of perceptual learning posited that the changes occurred in the earliest encodings because these areas have the same specificity – e.g. position and orientation – that has been reported in the perceptual learning literature. Petrov, et al, point out, however, that specificity in perceptual learning only requires that some part of the neural system responsible for making a particular decision have specificity, not that the changes that drive perceptual learning occur in units that have specificity. The solution, they argue, again involves selection. Learning can occur via changes in higher level structures and still have specificity, provided that those structures are reading the outputs of lower-level units that are specific to visual field location, orientation, or some other stimulus attribute. Specificity can arise from differentially selecting information from units with specificity, and can therefore arise from changes in higher-level, abstract representations of the relevant stimuli.

In order to unambiguously demonstrate lack of transfer, many perceptual learning experiments utilize two conditions (for learning and transfer testing) that involve distinct neural structures, e.g. stimuli at two orientations orthogonal to one another. In these experiments, lack of transfer is defined as training in one condition that does not enhance performance (or enhances performance less) in the other condition. This setup, Petrov, et al., argue, is poorly suited for discriminating between the receptive field modification hypothesis and the selective reweighting hypothesis. In this type of experiment, training in each condition involves distinct neural representations at the lowest level *and* distinct connections between these representations and higher level "decision units". Since both the representations and the connections are distinct, specificity could result from modification at either stage, and therefore cannot distinguish between the receptive field hypothesis (changes at the level of representation) and the selective reweighting hypothesis (changes in the connections).

In the Petrov, et al. experiments, instead of having a different stimulus in each condition, the same stimulus was presented in a different context in each condition.. In this case the representation at the lowest level is the same (i.e. the same units in V1), and the connection from these units to higher level units could be shared or distinct. Their experimental data were well-described by a model with a single representation of the stimulus, with training effects occurring in distinct connections between the representation and the decision units.

This recent work adds to earlier results that seemed discrepant with hypotheses of receptive field change. Ahissar, Laiwand, Kozminsky, & Hochstein (1998) trained subjects in a pop-out visual search task with a target of one orientation and distractors in another orientation. After learning had occurred, they swapped target and distractor orientations and again trained subjects to criterion. Finally, they switched back to the original orientations and found that performance was neither better nor worse than at the end of the first session of training. If training had caused modification of receptive fields of orientation-sensitive units in V1, they argued, then switching target and distractor orientations after training should have interfered with the earlier learning. Yet the earlier learning was preserved, and it

coexisted with the later learning. These results are more consistent with a model in which a particular task leads to selective use of particular analyzers, but the underlying analyzers themselves remain unchanged. This kind of result also addresses a deep concern about the idea of receptive field modification in vision: Given that the earliest cells in the striate pathway (beginning with cortical area V1) are the inputs for many visual functions, including contour, texture, object, space, and motion perception, task-specific learning that altered receptive fields at this level would be expected to affect or compromise many other visual functions. The fact that such effects do not seem to occur is comforting for visual health but less so for explanations of perceptual learning in terms of receptive field change.

The notion of selective reweighting corresponds nicely to Gibson's earlier account of perceptual learning as selection and the learning of distinguishing features. It raises the fascinating possibility that perceptual learning at all levels can be modeled as selection. One difference is that Gibson talked about selection of relevant stimulus information, whereas Petrov et al (2005) and others cast the selection notion in terms of selection and weighting of analyzers within perceptual systems. We believe these two versions of selection are two sides of the same coin. To be used, information must be encoded, and the function of encoding processes is to obtain information. Perceptual learning processes, we might say, select information relevant to a task by weighting heavily the analyzers that encode it. (Trying to distinguish more deeply between selection of information and selection of analyzers is much like puzzling over whether we really "see" objects in the world, or we see electrical signals in our brains. For discussion, including a claim that we see the world, not cortical signals, see Kellman & Arterberry, 1998, p. 10–11).

Relational Recoding

Neither receptive field modification nor simple selection from concrete inputs can account for the substantial part of human perceptual learning that is abstract. Take the idea of learning what a square is. A given square, having a specific size and viewed on a particular part of the retina, activates a number of oriented units in visual cortex. Suppose the learner is given category information that this pattern is a square. Using well-known connectionist learning concepts, this feedback could be used to strengthen the connection of each activated oriented unit with the label "square." One result of this approach is that with sufficient learning (to stabilize weights in the network), the network would come to accurately classify this instance if it recurred. Moreover, with appropriate design, the network would also be able to respond "square" to a display that contained a large part, but not all, of the previously activated units. These kinds of learned classification and generalization results are readily obtainable by well-established methods common in machine learning.

There is another level to the problem, however, and as we noted earlier, it is one with deep roots. At the turn of the last century, Gestalt psychologists criticized their structuralist predecessors for the idea that complex perceptions, such as perception of an object, are obtained by associating together local sensory elements. Their famous axiom "The whole is different from the sum of the parts" was aptly illustrated in many domains that are commonplace for perceivers but profoundly puzzling for the structuralist aggregating of sensory elements or, in our era, the weighting of concrete detector inputs. A square, for

example, is not the sum of the activations triggered by any particular square. Another square may be smaller, rotated, or displaced on the eye. All of these transformations activate different populations of oriented units in early cortical areas. Moreover, we also readily recognize squares made of little green dots, or squares made of sequences of lines that are oriented obliquely along the edges of a square. A square made in this way may have no elementary units in common with an earlier example of square, whereas another pattern that shares 90% of the activated units with an earlier example may simply not be a square at all. Being a square has to do with relations among element positions, not the elements themselves.

Kellman, Burke & Hummel (1999) made a proposal about how such abstract invariants may be learned. A crucial part of their approach is that the model be generic, in that it contains only properties we might expect on other grounds to be available in visual processing. That is, one could easily build a specific device to learn about squares, but the goal is to understand how we can be a general purpose device that learns the invariants of squares, rhombuses, dogs, cats, chess moves, etc. The key, according to these investigators, is a stage of early relational recoding of inputs. In vision, there are basic features and dimensions to which we are sensitive. Besides these, there must be operators that compute relations among features or values on dimensions. In the simulations of Kellman et al (1999), the key operator was an "equal / different" operator. This operator compared values, and produced an output that was large when two values were approximately equal and small when they differed beyond some threshold value. Their model learned to classify, from only one or two examples, various quadrilaterals (e.g., square, rectangle, parallelogram, trapezoid, rhombus, etc.). The model was given the ability to extract contour junctions and distances between these. The layer of equal / different operators scanned all pairs of inter-junction distances. A response layer that used the results of all equal / different operators readily learned the quadrilateral classifications from simple category feedback. Learning, even with a minimum of examples of each form class, generalized to patterns of different size, orientation, and position; with certain types of preprocessing, they would also generalize to figures made of any elements.

The key idea here is not to provide a complete account of learning in this particular domain but to begin to address the perceptual learning ability to discover higher-order invariants. We believe that along with the basic encodings of features and dimensions there is a set of operators that computes new relations. There is evidence for equal / different operators and one or two others. Some of these relational computations are automatic or nearly so; some may be generated in a search for structure over longer learning periods in complex tasks. The outputs of automatically generated relational recodings or newly synthesized ones that were not initially obvious may account for the salience of some relations in perception and for the longer term learning of other, less salient ones.

It is implied by this view that there is essentially a "grammar" of perceptual learning and learnability. Some relations in stimuli are obvious; some may be discovered with experience and effort; and some are unlearnable. These outcomes depend on the operators that compute relations on the concrete inputs, as well as on search processes that explore the space of potentially computable operations for classifications that are not initially obvious. If this

overall approach has merit, then determining the operators in this grammar, although difficult, stands as a key question in understanding abstract perceptual learning. The solution of this "Kantian" question of what relational operators we possess that may allow us to discover new invariants will help us understand how we learn to classify dogs and cats, become experts at chess and radiology, and will also empower artificial learning devices that may someday share in the remarkable levels of structural intuition displayed by human experts.

Abstract Relations and Language

Earlier we raised the question of whether human extraction of abstract relations requires capabilities extending beyond perception, such as symbolic processes in general or language in particular. We now return to this question in the context of modeling perceptual learning, as the answer we suggests fits within a general scheme of the process of discovery in perceptual learning.

Although we do not have a detailed process model of high-level perceptual learning, there are certain elements that seem likely to be part of any successful model. To see what these are, we focus on tasks in which the learner must discover what features or patterns appearing in particular instances (objects, events, situations) determine whether or not a particular instance is a member of a category. Such learning is common (as when a child learns the notion of dog). We assume for simplicity that the training consists of instances along with category feedback about whether particular instances are or are not in the category. We also assume that explicit instruction about which variables determine category membership is not given.

Any perceptual discovery task of this type involves what may be described broadly as a *search process*. The learner must somehow determine which attributes of the observed instances determine membership in a category. "Search" does not imply a particular mechanism. It is convenient to think of it as a process that involves sampling from a set of candidate features and then comparing the values of these features with outcome feedback. This sort of search can be implemented through gradual strengthening of various weights in a neural network, but not all inputs are available for adjustment on each learning trial. In human and animal learning, evidence suggests a sampling process involving limitations of attention (e.g., Trabasso & Bower, 1969), a point to which we return below. Our immediate focus, however, is on the potential candidate attributes that may determine a classification. Most obviously, these include perceptual features and relations that are naturally encoded. In classic discrimination learning work (some of which is summarized in Trabasso & Bower, 1969), humans or animals were tested for acquisition of correct responding where some outcome depended on highly salient stimulus features. In a typical experiment, figures (on a card or on a door in a jumping stand for rats) might be triangles vs. circles, red vs. black, and large vs. small. Learning involved determining which of these dimensions were relevant (e.g., color), as well as the significance of particular values on the dimension (e.g., red produces the reward). In cases where the perceptual dimensions are highly salient, the discovery problem in perceptual learning is minimized. With more complicated stimuli, it may be much harder to notice immediately what features and relations should be on the

candidate list. In fact, the candidate list for features or relations that could be crucial to some new classification seems unlikely to be a fixed list. In complex learning tasks, certain properties or relations that are initially unnoticed can be discovered. Advanced competence in chess, chemistry, radiology, mathematics, and many other domains almost certainly involves sensitivity to relations to which the novice is essentially blind. Earlier we discussed how new relations might be synthesized through a combination of initial perceptual encodings and the application of relational operators. This function – generating candidates for search -- seems to be required to explain advanced human perceptual learning. This view has similarities to the approach described by Perlovsky (2006) in that cognitive (and linguistic) models evolve to best correspond to and account for input signals.

In terms of process, the notion of search suggests that not all candidate features and relations are checked at the same time. This fact is consistent with the gradual nature of perceptual learning in many contexts and also with discrimination learning in simple contexts (Trabasso & Bower, 1969). Candidate features derived from basic perception and from the generation of higher-level relations must be checked for their correspondence with the categorization outcomes. Again, this description does not imply any particular implementation, but it does imply that not all candidates are checked at the same time. Such a limitation need not even arise specifically in the "checking" or weighting process; it might occur in capacity limits on what can be encoded in memory from perceiving each instance. Feedback based on the progression of learning could be used to determine whether to broaden, narrow or otherwise guide sampling. Obviously, many details of such a framework need to be further specified, and some suggest empirical tests.

For current purposes, however, the key idea is that of a ***salience hierarchy*** – some preferred ordering of candidate features. The idea that some properties are obvious at first glance and others require more effort to notice seems straightforward: it is required by the fact that some perceptual discovery tasks are trivially easy whereas others require extended learning. This aspect of perceptual learning is the likely locus for effects of symbolic learning and language. By assigning symbols to particular properties or relations, as by naming a relation, its salience may be increased. Increased salience for some property leads to a greater likelihood that it will be searched for, encoded, and checked for relevance to the categorization outcomes.

Humans may well exceed animals in their extraction of abstract relations from objects, arrangements and events, and this difference may depend on language. The dependence, however, does not imply a fundamentally different process from perceptual learning as we have described it. Rather, the effects of language and symbolic experience may influence salience hierarchies in search processes involved in perceptual discovery. As with other aspects of perceptual learning, these effects may often be highly task specific: prior experience with and symbolic representation of a property may lead to greater likelihood of encoding that property or selecting it to check its relevance to a classification being learned. Although these conjectures are obviously ripe for further research, existing evidence seems consistent with them. Studies of concept formation in pigeons (e.g., Huber, 1999) do not indicate a clear limit on pickup of abstract relations. They do indicate that pigeons extract relations less readily than do humans but that under optimal conditions, pigeons seem to

classify based on abstract relations. These observations are consistent with the idea that making certain relations salient is helpful in abstract perceptual learning; in humans, language may contribute to making a relation salient.

It is also possible that language may allow relations that are recognized in tasks in one cognitive domain to be selected in another domain. Spelke (2003) has suggested that language can serve this kind of integrative or connective function: Core cognitive domains may develop prior to language, but their representations are often narrowly confined. Language may facilitate connections across domains. In guiding perceptual learning of high-level relations, language may point the way to extraction of relations in one domain because these have been relevant and symbolically represented in another. Such phenomena that appear to involve abstract learning differences between human children and adults (e.g., Kendler & Kendler, 1961), as well as the observation that some relational learning tasks can be acquired only by humans and symbol-trained chimpanzees are consistent with this role of language.

Presupposed by these interpretations, however, is the idea that the relevant relations can be found in objects and events if made salient. In other words, the capability to encode relevant properties needs to be present. Such capabilities may indeed differ across species, such that in some species, no amount of directing of attention could lead to discovery of a certain relation. If evolution has overseen some progression in perceptual abilities to pick up abstract relations, that would certainly be relevant to understanding differences in intelligence across species. Clearer understanding of species' differences in this regard would be very useful in considering this idea. Moreover, there is a fascinating chicken-and-egg problem regarding symbolic language and abstract perceptual encoding: It is possible that evolutionary development of sensitivity to abstract stimulus relations contributed to the emergence of language. Although there are obviously many potential influences on the evolution of language, at minimum, some abstract perceptual learning capabilities are crucial. As we consider below in Section VI, a child's ability to classify parts of the speech stream into grammatical categories involves complex cases of abstract perceptual learning, and these appear to have innate foundations (Hirsh-Pasek & Golinkoff, 1997).

V. The Scope of Perceptual Learning

It is apparent from our discussion that perceptual learning encompasses several information processing changes. These are linked by our general definition of perceptual learning as comprising improvement in the pickup of information. There appear to be different ways to improve: coming to select relevant features for some classification, discovering higher-order invariants to which the perceiver is initially insensitive, and becoming more fluent or automatic in information pickup.

It would be possible to classify these effects as deriving from different learning processes. As research progresses, it may turn out that there are a variety of mechanisms at work in these changes. Also, some characteristics of perceptual learning may be shared with other kinds of information processing. For instance, information extraction may become more automatic with practice, but so might familiar patterns of reasoning or inference.

Still, there are overriding functional reasons for considering perceptual learning, broadly construed, as a coherent category in a taxonomy of learning. Not only do the perceptual learning changes we have considered involve improvement in the way information is acquired and occur under the same conditions, but they also appear to be deeply interconnected. Discovery and fluency effects, as Bryan & Harter illustrated long ago, work in synergistic and recursive fashion to produce a spiral of access to progressively higher-level structure. This spiral of discovery → fluency → new discovery is likely responsible for many of the most impressive of human accomplishments. Also common to the family of perceptual learning effects is that they are not easily explained by other common notions in the taxonomy of learning.

Grouping phenomena together makes most sense if doing so helps to reveal general underlying principles. We believe this is the case with many perceptual learning phenomena, and in what follows we offer some analysis relevant to both that goal and to some current views about perceptual learning that suggest a much narrower scope.

The Current Modal View of Perceptual Learning

As we have considered, the recent surge of interest in perceptual learning has focused almost exclusively on basic sensory acuities. Many experiments reveal to the observer two stimulus displays, with instructions to the observer to judge, over thousands of trials, which of the two has been presented each time.

Such a focus may seem puzzling in terms of perceptual learning's role in real-life learning tasks, few if any of which resemble this scheme. It is also historically aberrant. In Gibson's classic (1969) review, there is consistent emphasis on extraction of invariance from variable instances but very few studies that conform to the more recent paradigm. The differences are not merely methodological, however; there is a contrast between the task content: shape, arrangement, and relational structure for many tasks in the earlier generation of research, contrasted with current work involving discrimination tasks such as Vernier acuity.

The recent focus is special but not coincidental. It rests on a plausible and coherent set of ideas. We will argue that the ideas are not correct, but their appeal is nonetheless easy to understand.

The primary organizing idea in recent work in perceptual learning is that perceptual learning can provide a window into plasticity in the brain. The sensory acuities found to be malleable through learning are also the functions currently best understood in terms of brain anatomy and physiology. For example, the first visual cortical area (V1) is known, from single-cell recording work in animals, to have a preponderance of cells sensitive to contrast at particular locations, orientations, and particular ranges of spatial frequencies (Hubel & Wiesel, 1968; DeValois & DeValois, 1988). The receptive fields for such units have been modeled as a Gabor filters – periodic (e.g., sinusoidal) functions of contrast, multiplied by a Gaussian window. Moreover, it is known that orientation sensitivity, in primates, is not present in this visual stream prior to V1. Thus, a perceptual learning task involving orientation sensitivity may use a pair of Gabor patches differing slightly in orientation, accompanied by an assumption that improvement in performance will reflect changes in mechanisms at this

early processing level. Indeed, some have argued that the definition of perceptual learning, or one especially important kind of perceptual learning, should be confined to changes at this level. For example, Fahle & Poggio (2002) defined perceptual learning as encompassing "parts of the learning process that are independent from conscious forms of learning and involve structural and/or functional changes in primary sensory cortices."

Consistent with these ideas are some common phenomena from perceptual learning tasks involving basic sensory dimensions. It is often found that learning effects in perceptual learning tasks show limited transfer across changes in the stimulus or the location of presentation (e.g., the eye, hemifield, or retinal location use during learning). Limited transfer has routinely been interpreted as showing a low-level locus of learning. For example, an effect found during training on a certain location in one eye that does not transfer to the same position in the visual field on the other eye has been taken to indicate that the locus of learning must be at a level prior to that where binocular combination of inputs occurs (e.g., V1).

The final assumption in this perspective is that behavioral results involving perceptual learning in animals, that seem analogous to human results, can be directly linked to receptive field changes in primary sensory cortices using direct physiological measures (e.g., Recanzone et al, 1993). Taken together, one can see the coherence of these basic ideas about perceptual learning and early cortical change: behavioral results that involve basic sensory discriminations; transfer data (often, failure to transfer) that suggest a low-level locus of learning, and direct physiological evidence of neural change from animal experiments.

Critique of the Modal View

Despite its plausibility, there are reasons to question the currently prevailing orthodoxy about perceptual learning. As research has progressed, it has cast doubt on each element of the set of assumptions that justifies a low-level sensory emphasis. Conversely, we believe that evidence has begun to reveal more general principles of perceptual learning that unify tasks and results at different levels. Perceptual learning in basic sensory tasks and higher level tasks may be more closely related than anticipated. Discovering and elaborating general principles of perceptual learning may require abandoning the narrowest views of perceptual learning.

Specificity of learning does not implicate low-level mechanisms.

Recall that specificity of learning, or equivalently, lack of transfer of learning across stimulus conditions has been traditionally interpreted as a sign that the neural changes associated with the learning occur earlier on in visual processing. This is simply because neurons in the early visual system are known to encode visual stimuli in a manner that is specific to, for example, a position on the retina. The lack of transfer in perceptual learning tasks is not, however, a particularly robust effect. Across learning domains and for studies within particular domains, transfer results have been inconsistent (for a review, see Kellman, 2002). More recently, Young, Li, Levi & Klein (2005) re-examined the idea that perceptual learning in a Vernier acuity task is specific to the eye trained. They pointed out that in previous studies, in which learning effects were found to be specific to the trained eye, the

observer's untrained eye was covered by a patch. This condition may have favored learning to ignore the patch, which could produce an interfering visual input. When they performed the same experiment with a light diffuser over an open untrained eye, they found complete transfer of learning effects to the untrained eye. This kind of result suggests that the specificity of perceptual learning effects may reflect the task the experimenter has chosen (sometimes in unsuspected ways) rather than constraints on learning mechanisms.

Other research confirms that small task variations can lead to large differences in generality of transfer. For example, Liu (1999) found substantial perceptual learning effects in discrimination tasks both for motion directions differing by 3 deg (as previously studied by Ball & Sekuler, 1982) and for motion directions differing by 8 deg. When he tested for transfer, however, learning of the 3 deg discrimination showed little transfer to new orientations, but learning of an 8 deg discrimination transferred robustly. In the face of findings like these, defining perceptual learning by specificity of transfer would have paradoxical consequences (i.e., learning to discriminate motion directions differing by 3 deg would be perceptual learning, but learning to discriminate motion directions differing by 8 deg would not be). Likewise, following the typical inferences from specificity of transfer, we would have to infer different cortical loci for these two very similar tasks. As we will suggest below, the variability in transfer results is consistent with selection models of perceptual learning, and such models avoid these counterintuitive consequences.

An attack on the typical logic of inferring a low-level locus from transfer was put forth by Mollon & Danilova (1996), who argued that the idea is simply a fallacy (c.f., Dosher & Lu, 1998). Specificity of transfer need not reflect a low-level site of modifications in the nervous system; it is equally consistent with a selection notion: that more central processes select which outputs of relevant analyzers at lower levels are relevant to a particular task. In this view, not only does learning involve higher levels, but no changes at lower levels are required, and specificity in transfer mostly reflects the task the experimenter has chosen. This view is also more consistent with the large variability in transfer outcomes with relatively small variations in procedures and tasks.

Receptive fields in early visual areas do not show sufficient changes to account for perceptual learning effects.

The argument for low-level loci of learning in humans has been based on the specificity of perceptual learning effects. In light of both the logical issues and inconsistent results described above, there is little or no evidence that any perceptual learning results in humans involve changes in primary sensory cortices. Such a statement will surely sound like heresy to many researchers who have presumed they are studying such changes. They would most likely respond by pointing to animal research in which behavioral tasks and physiological changes can both be assessed. Some evidence has suggested such changes due to perceptual learning in primary auditory and somatosensory cortex. Yet, in vision, there is little evidence from animal research for modification of early areas; changes in receptive field structure or dynamics following perceptual learning have not been consistently observed (e.g. Crist, Li, & Gilbert, 2001; Schoups, Vogels, Qian, & Orban, 2001). Training that results in large behavioral changes, in terms of, e.g., discrimination improvement, may occur despite little

corresponding changes in the receptive fields of V1 and V2 cells (in monkey). Where changes have been found, the results have been subtle, and of too small a magnitude to account for the changes to behavior (Ghose, Yang, & Maunsell, 2002). More robust training-induced modifications have been found at later stages of processing, e.g., cells in area V4 (Yang & Maunsell, 2004). The greater complexity of the receptive fields of these cells (Pasupathy & Connor, 2001) makes characterization of the impact of these changes challenging. In any case, these changes are well beyond "primary sensory cortices," and the observed changes in V4 could be consistent with the neural implementation of either receptive field change or selective reweighting models of perceptual learning.

Higher level variables are involved in perceptual learning.

A variety of findings indicate that, even in low-level tasks, higher level variables affect perceptual learning. Attention to particular stimulus attributes, having a task, and/or an internal reward signal upon finding a target seems important (Seitz & Watanabe, 2003; Shiu & Pashler, 1992), although under certain circumstances these variables may allow learning to occur for spatially and temporally coincident but task-irrelevant stimuli (see Seitz and Watanabe, 2005).

Ahissar & Hochstein (2004) suggested that perceptual learning always involves an interplay of higher and lower levels of processing. In their "reverse hierarchy" approach, learning is initiated through high level attention and task involvement. Learning descends to the lowest processing level at which task-relevant regularities exist. This approach thus involves discovery or selection, not just of specific analyzers at some level, but among possible levels of information and processing as well.

Perceptual Learning should not omit perception.

Although it is common to label early cortical changes as perceptual learning, and some have suggested or assumed that only such changes qualify as perceptual learning, this view generates a seldom-noticed paradox. Learning effects, if confined to primary sensory cortices, would likely not be *perceptual* learning. In vision, the earliest cortical areas (V1, V2) contain cells with receptive fields sensitive to oriented contrast in small regions, in specific spatial frequency ranges. These areas are unlikely to contain representations of what objects are in a scene, their shapes, sizes, or the 3D arrangements of objects and surfaces. Even a single piece of an abrupt luminance edge of an object against a background cannot be represented by any one such detector. Most perceptual descriptions are likely computed further along in the visual pathway, such as object shape in parietal regions and LOC (Kourtzi & Kanwisher, 2001) or spatial arrangement obtained from a synthesis of varying spatial cues, which may be registered in the caudal intraparietal sulcus (Tsutsui, Sakata, Naganuma, & Taira, 2002). Even a simple property like perceived object size is unlikely to be coded by cells in V1 or V2, whose responses, unlike perceived size, will change as the same object is viewed from different distances. We are only beginning to understand how the outputs of simple and complex cells in early cortical areas are transformed into representations of contours and surfaces with properties like color and transparency (e.g. Grossberg, 1997; see Neumann, Yazdanbakhsh, & Mingolla (2007) for a review of computational modeling efforts in this area). Perceptual learning effects important in real-

world tasks (such as shape discrimination) often involve these higher-level properties (see Grossberg, 1999 for a discussion of computational modeling approaches to understanding learning in cerebral cortex).

Perceptual learning   sensory plasticity.

Perhaps the problem in a notion of perceptual learning that omits most of perception is that we need to distinguish notions of perceptual learning and sensory plasticity. The issue of when and how cortical receptive fields can change is certainly an important topic in its own right. Such changes may occur in some perceptual learning tasks, and they may occur during early visual development, or due to deterioration when stimulation is excluded, and in other circumstances.

Yet one must wonder why changes, if confined to the earliest cortical areas, should be considered perceptual learning at all. Indeed, recent findings (Garrigan & Kellman, 2008) give reason for believing that the perceptual learning may actually operate only on interpreted perceptual representations. In our current understanding of visual processing, there is little reason to suppose that most, if any, such representations reside in the earliest cortical areas. Perceptual learning involving interpreted perceptual representations may well produce changes at the earliest sensory levels, but learning effects would unlikely be *confined* to such plasticity. There are also many important questions of sensory plasticity involving these early areas that do not necessarily involve learning effects, such as those relating to maturation and tuning of cortical units in infancy (Norcia & Tyler, 1985). We believe that perceptual learning and sensory plasticity label two important notions; they may often be related but they are not interchangeable.

Selection as a Unifying Principle in Perceptual Learning

In addition to the foregoing arguments against confining perceptual learning to sensory plasticity at the earliest cortical levels, it should already be clear that there is an important theoretical consideration against doing so. Earlier we described the case for modeling even low-level perceptual learning as selection. In her classic work on perceptual learning, Gibson saw selection as a unifying principle in higher level as well as low-level perceptual learning contexts. Whether one sees as more important the window into cortical plasticity opened by low-level perceptual learning or the role of perceptual learning in richer, real-world learning domains, restricting the label "perceptual learning" to one part of the landscape would miss an obvious general principle that unifies perceptual learning as a type of learning. The specific information and analyzers involved may no doubt be important for understanding particular cases of perceptual learning, but the overarching explanation of perceptual learning as involving task-driven discovery of relevant information for classifications should not be missed.

Receptive Field Change and Selection

A related insight is that even receptive field changes that could occur after learning may ultimately require a selection explanation as well. Imagine a cortical cell that becomes sensitive to a larger retinal area as a result of training. A more detailed account of this change, if it were to become available, would no doubt show that this cortical cell's links

with retinal ganglion cells had changed. Ganglion cells gathering information from spatial regions that previously did not affect the cortical cell now would affect it. The changing of connections or weights with earlier processing levels is again a notion of selection or selective reweighting. Thus, even cases of receptive field change may involve selection, if viewed more finely. This does not imply, however, that all selection may be thought of as receptive field change. There are certainly levels in the nervous system at which decision processes utilize information from various inputs yet we would not be inclined to characterize these as having receptive fields in anything like the usual sense. Such decision processes may not be realized in terms of outputs of single neural units, and their inputs would not be maps on the receptive surfaces of sensory systems. More likely, decision processes at different levels can gain access to outputs of various perceptual processes, and the improvement in selectivity and fluency that emerges from practice need not invoke receptive field concepts.

Another important consideration is that receptive fields are relatively permanent (which is why the hypothesis of receptive field change through learning is so interesting in the first place). In contrast, reliance by decision processes on particular sets of analyzers may change from task to task. This impermanence of task-specific selection has been the key finding in some empirical findings that have led investigators to argue for selection models rather than receptive field change (Ahissar et al, 1998; Petrov, et al, 2005). In some cases, especially involving the lowest levels and perhaps particularly for senses other than vision (given current evidence), improvements in selectivity may well occur through receptive field change. It is a question of continuing interest if, when, and where, this occurs.

Perceptual Learning and Perceptual Constancy

Reframing perceptual learning in terms of selection, with access not restricted to the simplest sensory analyzers, raises an interesting question. What, if any, are the constraints on selection? Is information at any level of processing – from the earliest analyzers to the most complex perceptual representations – available to perceptual learning processes? As we have seen, some have suggested that only activity of analyzers at the earliest levels falls within the domain of perceptual learning. Paradoxically, it may actually be the case that the earliest analyzers are never the direct targets of learning processes. These ideas have recently emerged in studies relating perceptual learning to perceptual constancy.

Perceptual constancy refers to the fact that perceptual representations (and experience) involve stable properties of objects, not the sensory inputs used to derive those representations. For example, the perceived size of a person does not change as the person walks toward you or away from you. Such events do change the size of the projection of the person on your retinas. In this case, size constancy describes the fact that you perceive an unchanging physical size despite changes in the projective (retinal) size. Under common conditions, the projective size is (along with registered distance) an input into the computation that produces size constancy. Constancy processes produce perceptual representations that correspond to enduring physical properties of objects, despite changes in the sensory inputs (the proximal stimuli) used to compute perceptual representations. A deep reason why constancy processes are important is that we perceive by means of stimuli

received at the senses (energy), but ecologically, what we need to know about (primarily) is the structure of material objects, spatial arrangements, and events (Garrigan & Kellman, 2008). Constancy processes enable us to deal with important and relatively stable material structures in the physical world which are made known to us by constantly fluctuating arrays of energy reaching our senses.

If constancy-based representations are important in our descriptions of the environment, might they also play a special role in learning? Garrigan & Kellman (2008) studied this question in series of experiments. Using established relations of perceptual constancy and sensory inputs (such as the use of retinal size in computing perceived size), they tested subjects in two conditions in each experiment. In each condition, subjects' task was to discover the basis of a classification. On each trial, they were shown displays and asked to indicate whether or not the display fit in category X. The relevant properties were not described to subjects; they were given accuracy feedback only, and their task was to discover any stimulus properties that could lead to accurate classification. In one condition, in each of the basic experiments, subjects were shown a display in which either a perceptual (constancy-based) invariant or a sensory invariant could in principle be discovered and used to do the classification. In the case of size perception, the stimulus on each trial consisted of a pair of rectangles that were the same in retinal size and also in perceived size (because they were at both at the same distance from the observer). If perceptual learning processes can discover invariants at any level of processing, learners could succeed in this task either by extracting the perceptual regularity or the sensory one (the relation of retinal sizes). In a separate condition, displays could be successfully classified only by using the relation of retinal size. This condition was arranged by using identical stimuli as in the first condition, but using stereoscopic disparity to vary the apparent depth of the two displays. The question was whether a sensory invariant could be discovered if it did not correlate with a perceptual invariant. This question was asked in three different perceptual domains – perception of size, surface lightness, and motion. In every case, learning readily occurred when a perceptual invariant was available to be discovered. When only a sensory invariant was available, there was no evidence of learning, even after hundreds of trials. The results were especially revealing in that in each case, the sensory invariant is known to be an ingredient used in computation of the perceptual invariant. The visual system picks up the sensory information, but that information does not seem accessible to perceptual learning processes. A separate control using perceptual equivalence that was uncorrelated with sensory equivalence also showed robust learning effects. (This control indicated that a perceptual regularity alone was sufficient for learning, not a combination of a perceptual and sensory regularity.) Figure 4 illustrates one of the experiments – a study of perceptual learning in surface lightness.

This result fits well with a model of perceptual learning that incorporates higher level, abstract representations, and suggests a strong constraint on perceptual learning. Perceptual learning is in fact perceptual. Information in early encodings that does not survive transformation to the percept is unavailable. This assessment inverts some popular wisdom. Far from being confined to the earliest analyzers, perceptual learning may not directly access early sensory analyzers at all.

It is important to note that this view does not imply that neural changes associated with perceptual learning cannot occur in early cortical regions. It simply means that adjustments that arise from learning throughout the entire system work through constancy-based representations, a point related to the claim of Ahissar & Hochstein (2004) that lower level changes in perceptual learning are guided by higher levels.

Ecologically, it makes sense that learning processes access perceptual encodings, not the sensory inputs used to compute them. Red objects in white light might reflect mostly red light onto the retina, but red objects in blue light might reflect mostly blue light onto the retina. Discovering that the apple is red requires considerable processing, but that persisting property of apples is likely to be the kind of information useful in learning important regularities about the environment as well as in guiding thought and action. Perceptual learning utilizes distinctions at these more abstract levels because these levels capture behaviorally-relevant information rather than incidental details of viewing conditions. Without such constraint, learners might swamped search processes that lead to discovery in perceptual learning. Search processes that need to explore a large space of possible relations among the activity of photoreceptors on the retina may be intractable; search processes that consider structural properties of objects may be both tractable and adaptive. Perceptual learning involves selection, and available evidence suggests that selection depends on constancy-based, perceptual representations.

VI. Perceptual Learning and Instruction

If one consults the educational literature about the cognitive bases of human learning, one would find ample treatment of fact and concept learning, conceptual understanding, procedure learning, constructing explanations, thinking and reasoning. Except for an occasional mention of pattern recognition, perceptual learning would be starkly absent. Yet it is arguably one of the most, possibly *the* most, important component of human expertise.

If one looks instead at research, not on education, but on expertise, perceptual learning stands out. DeGroot (1965), himself a chess master, studied chess players, with the expectation that master level players considered more possible moves and countermoves or in some sense thought more deeply about strategy. Instead, he found that their superiority was shown mostly on the perceptual side. They had become able to extract meaningful patterns in larger chunks, with greater speed and less effort than less skilled players. De Groot suggested that this profile is a hallmark of human expertise in many domains:

> We know that increasing experience and knowledge in a specific field (chess, for instance) has the effect that things (properties, etc.) which, at earlier stages, had to be abstracted, or even inferred are apt to be immediately perceived at later stages. To a rather large extent, abstraction is replaced by perception, but we do not know much about how this works, nor where the borderline lies. (de Groot, 1965, pp. 33–34)

Such differences between experts and novices have since been found to be crucial in research on expertise in a variety of domains, such as science problem solving (Chi et al, 1981; Simon, 2001), radiology (Kundel & Nodine, 1975; Lesgold, et al., 1988), electronics

(Egan and Schwartz, 1979), and mathematics (Robinson & Hayes, 1978). An influential review of learning and its relation to education (Bransford, Brown & Cocking, 2001) summed it up this way:

> Experts are not simply "general problem solvers" who have learned a set of strategies that operate across all domains. The fact that experts are more likely than novices to recognize meaningful patterns of information applies in all domains, whether chess, electronics, mathematics, or classroom teaching. In deGroot's (1965) words, a "given" problem situation is not really a given. Because of their ability to see patterns of meaningful information, experts begin problem solving at "a higher place" (deGroot, 1965). (Bransford, Brown & Cocking, 1999, pp. 48)

As was evident in our discussion of specific examples, perceptual learning serves in the development of expertise in multiple ways, such as automaticity and chunking in information pickup. Table 1 summarizes these changes, all of which have been demonstrated in studies of expert performance and experiments on perceptual learning.

Looking over the effects in Table 1, one is struck by the remoteness of these characteristics that drive expert performance from the usual products and goals of school (and university) learning. It is not the case that the facts, concepts, and procedures taught in school are unimportant. It is more that those foci alone leave out something important. How does hearing a lecture help a university student to do more parallel processing or pick up larger chunks of information at a glance? The contrast between the usual emphases in conventional instruction and the characteristic strengths of experts are relevant to understanding the potential role perceptual learning could play in instructional contexts and how it might best mesh with and enhance conventional methods.

Natural Kind Learning

Moving from experts to the other side of the continuum, it is interesting to consider that before a child goes to school, discovery processes in perceptual learning, including highly abstract ones, are very much in evidence. Imagine a young child going for a walk with her father. Upon seeing a dog, the child points, and the father says "That's a dog." Suppose this particular dog is a small, white, poodle. On some other day, the child sees another dog -- this one a large, black, Labrador retriever. Again, someone says "dog." And so on. With each instance, something about a particular dog (along with the label "dog") is encoded. As this process continues, and a number of instances (probably not a particularly large number) have been encountered, the child becomes able to look at a new, never before seen dog and say "dog." This is the magical part, as each new dog will differ in various ways from any of the examples encountered earlier. Moreover, with practice, the child will correctly distinguish novel instances of dog, cat, squirrel, etc. from each other. A particular cat or squirrel may have properties that resemble some known dog; a black dog and a black cat are more similar in color than are a black and white dog. Despite similarities of instances in different classes and differences within classes, the learner somehow comes to extract properties that can be used to classify novel instances accurately. Such examples appear to be cases of abstract perceptual learning, as the most crucial information allowing assignment to the category very likely involves relational variables rather than the most concrete

features, such as color. A miniature, purple, plastic model of a dog wearing tennis shoes would still be recognized as a dog.

These properties underlying a classification can be highly complex and implicit. If a child, or even an adult, is asked to state a set of rules that would allow a novice to distinguish dogs, cats, and wolves, they cannot ordinarily do so. In contrast with much of school learning to come, it is striking that these feats of natural learning occur without the child ever being presented with a lecture on the distinguishing features of dogs or cats. Rather, structure is extracted from instances and category feedback. Such perceptual learning processes are crucial not only for developing understanding of the objects and events in the world; they also play a pivotal role in language acquisition, at multiple levels. Concepts like noun, verb, adverb, and preposition are taxing enough when taught explicitly in middle school. How is it that these abstract classes are extracted and used in language acquisition, allowing grammatical structures to be processed (e.g., Hirsh-Pasek & Golinkoff, 1997) and facilitating the learning of new words? At a different level, learning may be involved in the ability of the young language learner to detect invariance in the structure of speech signals across different speakers. Evidence suggests that the perceptual learning processes needed for these achievements, including clear cases of abstract perceptual learning, are present relatively early in infancy (Gomez & Gerken, 2000: Marcus et al, 1999; Saffran, Loman, & Robertson, 2000).

The Container Metaphor

Once in school, learning will be organized and guided by certain deep-seated assumptions about what learning is. Both formal and informal discussions of learning, both in educational and commercial training settings, most often consists of two categories: declarative and procedural learning. The former comprises facts and concepts that can be verbalized and the latter of sets of steps that can be executed. Bereiter & Scardamalia (1998) suggested that this pervasive, implicit view is a "folk psychology" notion of learning. They call it the *container metaphor*, summed up by saying that "Knowledge is most readily conceived of as specifiable objects in the mind, such as discrete facts, beliefs, ideas…" and learning "… involves retaining and retrieving such objects." (Bereiter & Scardamalia, 1998, p. 487). As Figure 6 shows, this view has also been shared by cartoonist Gary Larson.

We do not mean to suggest that perceptual learning is purposely neglected. Indeed, as we have suggested, one problem is that both learning scientists and educators have had little acquaintance with perceptual learning. The identification of teaching, learning, and knowing with the explicit is deep-seated. Facts and concepts that can be stated and explicit procedures – sequences of operations that can be enacted – are the mainstay of most activities we call educational. We believe the lack of focus on perceptual learning also derives from a lack of suitable techniques. How might we teach pattern recognition or structural intuition? Some perceptual learning surely occurs by considering examples in lectures or homework, but these activities target perceptual learning obliquely at best. It is often recognized that there are aspects of learning that extend beyond the classroom – the novice pilot, radiologist, chemist, or computer programmer may be told that fluency and the intuitive grasp of structure will come from practice or seasoning. A special obstacle for any other approach to

these aspects of learning is that they often involve unconscious processing. If not even the skilled expert can verbalize the information being used to do a task, how can a teacher convey this to students? Even when details can be expressed, hearing a description does not make the student a fluent extractor of information.

Technologies for Perceptual Learning

Research in recent years suggests that there are ways to address perceptual learning systematically in instructional settings. These can accelerate the growth of perceptual expertise. Indeed, the reader may have already noted the irony of educators lacking methods to address perceptual learning while researchers produce perceptual learning routinely in almost any laboratory task they select. Of course, there are differences between most laboratory tasks and richer educational domains, but equally surely there is significant potential to address perceptual learning in instruction.

A number of recent efforts have produced markedly successful outcomes using perceptual learning techniques. Kellman & Kaiser (1994) applied perceptual learning methods to pilot training, using what they called Perceptual Learning Modules (PLMs). In a Visual Navigation PLM, pilots learned navigational skills by mapping, on short, speeded trials, videotaped out-of-the-cockpit views onto locations shown on standard visual navigation (VFR sectional) charts. Remarkable improvements in accuracy and speed occurred in just an hour of training, even among experienced aviators. In a separate study of flight instrument interpretation, pilots classified aircraft attitude (e.g., climbing, turning) from an array of primary flight displays used by pilots to fly in instrument conditions. They found that under an hour of training allowed novices to process configurations more quickly and just as accurately as civil aviators who had on average 1000 hours of flight time (but who had not used the PLM). When experienced pilots used the PLM, they also showed substantial gains, paring 60% off the time needed to interpret instrument configurations.

Perceptual learning interventions to address speech and language difficulties have been reported to produce benefits (Merzenich, et al., 1996; Tallal, Merzenich, Miller, & Jenkins, 1998). For example, Tallal et al. showed that auditory discrimination training in language learning using specially enhanced and extended speech signals improved both auditory discrimination performance and speech comprehension in language-impaired children.

Applications in medical and surgical training illustrate the value of perceptual learning in addressing dimensions of learning not encompassed by ordinary instruction. Guerlain et al (2004) applied PLM concepts to address issues of anatomic recognition in laparoscopic procedures. They found that a computer-based PLM approach patterned after the work of Kellman & Kaiser (1994) led to better performance than traditional approaches. The training group presented with variation in instances selected to encourage learning of underlying invariance later showed improvement on perceptual and procedural measures, whereas a control group who saw similar displays but without the structured PLM did not. Their data implicated perceptual learning as the source of the improvement, as neither group advanced on strategic or declarative knowledge tests.

Perceptual learning technology is also being applied to high-level educational domains such as mathematics and science. Although these subjects involve a variety of cognitive processes, they rely substantially on pattern recognition and fluent processing of structure, as well as mapping across transformations (e.g., in algebra) and across multiple representations (e.g, graphs and equations). These aspects are not well-addressed by conventional instruction, and a variety of indicators suggest that they may be disproportionately responsible for students' difficulties in learning (Kellman et al, in press). Although this research area is relatively new, findings indicate that even short PLM interventions can accelerate fluent use of structure, in contexts such as the mapping between graphs and equations (Silva & Kellman, 1999), apprehending molecular structure in chemistry (Wise et al., 2000), processing algebraic transformations, and understanding fractions and proportional reasoning (Kellman et al, in press).

Although a full description is beyond our scope here, it may be helpful to mention some elements of perceptual learning interventions. A key assumption is that information extraction advances when the learner makes classifications and (in most cases) receives feedback. Digital technology makes possible many short trials and appropriate variation in short periods of time, allowing the potential to accelerate perceptual learning relative to less frequent or systematic exposure to structures in a domain. Unlike conventional practice in solving problems, learners in PLMs discriminate patterns, compare structures, make classifications, or map structure across representations. In an Algebraic Transformations PLM, for example, learners saw a target equation on each trial and made a speeded choice indicating which of several alternative equations was a legal transformation of the target (Kellman et al, in press). Although learners never practiced solving equations in this learning intervention, two to three 40-minute sessions of this sort led to strong advances in fluency in algebra problem solving (specifically, an average decrease from about 28 sec to about 12 sec per problem). Another key element of PLMs is the novelty of instances. Learning to process structure that defines a category, as opposed to labels for particular memorized instances, requires that most learning items (and transfer items) be novel. Here again, digital technology – making possible the storage or generation of large numbers of novel instances – is an especially good fit for this approach to learning.

There are many other issues in structuring learning in PLMs – issues of sequencing, feedback, variation of positive and negative instances of categories, mixing of learning tasks, integration of perceptual learning activities with conventional instruction, and so on. Moreover, significant research remains to be done regarding how to optimize these interventions and combine them with instruction. It is already clear, however, that perceptual learning interventions offer great promise in addressing neglected dimensions of learning and overcoming common difficulties in learning. If these possibilities are realized, they will comprise a far-reaching consequence of basic scientific efforts to bring to light and understand perceptual learning, both as an important variety of learning and as one that interacts with and supports other forms of learning and cognition.

Acknowledgments:

References

Ahissar M & Hochstein S (2000). Learning pop-out detection: The Spread of Attention and Learning in Feature Search: Effects of Target Distribution and Task Difficulty. Vision Research, 40, 1349–1364. [PubMed: 10788645]

Ahissar M, & Hochstein S (2004). The reverse hierarchy theory of visual perceptual learning. Trends in Cognitive Sciences, 8(10), 457–464. [PubMed: 15450510]

Ahissar M, Laiwand R, Kozminsky G, & Hochstein S (1998). Learning pop-out detection: Building representations for conflicting target-distractor relationships. Vision Research, 38(20), 3095–3107. [PubMed: 9893818]

Anderson JR, Corbett AT, Fincham JM, Hoffman D & Pelletier R (1992). General principles for an intelligent tutoring architecture. In Regian J & Shute V (Eds.), Cognitive approaches to automated instruction (pp. 81–106). Hillsdale, NJ: Erlbaum.

Ball K, & Sekuler R (1982). A specific and enduring improvement in visual motion discrimination. Science, 218(4573), 697–698. [PubMed: 7134968]

Baron J (1978). The word-superiority effect: perceptual learning from reading. In Estes WK (Ed.), Handbook of learning and cognitive processes Vol. 6 Hillsdale, NJ: Erlbaum.

Bereiter C & Scardamalia M (1998). Rethinking learning. in Olson D & Torrance N (Eds), The handboook of education and human development (pp 485–512). Malden, MA: Blackwell.

Berkeley G (1709/1910). Essay towards a new theory of vision London: Dutton.

Bransford JD, Brown AL, & Cocking RR (Eds). (1999). How people learn: Brain, mind, experience, and school Washington, D.C.: National Academies Press.

Brunswik E (1956). Perception and the representative design of psychological experiments Berkeley: University of California Press.

Bryan WL and Harter N (1897). Studies in the physiology and psychology of the telegraphic language. Psychological Review, 4, 27–53.

Bryan WL & Harter N (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. Psychological Review, 6(4), 345–375.

Bushnell IWR, Sai F & Mullin JT (1989). Neonatal recognition of the mother's face. British Journal of Developmental Psychology, 7, 3–15.

Chase WG & Simon HA (1973). Perception in chess. Cognitive Psychology,. 4(1), 55–81.

Chi MTH, Feltovich PJ, and Glaser R 1981 Categorization and representation of physics problems by experts and novices. Cognitive Science 5:121–152.

Clawson DM, Healy AF, Ericsson KA, and Bourne LE (2001) Retention and transfer of morse code reception skill by novices: Part-whole training. Journal of Experimental Psychology: Applied 7(2), 129–142. [PubMed: 11477980]

Cover TM and Thomas JA (1991). Elements of Information Theory New York: Wiley.

Crist RE, Li W & Gilbert CD (2001). Learning to see: experience and attention in primary visual cortex. Nature Neuroscience, 4(5), 519–525. [PubMed: 11319561]

de Groot AD (1965). Thought and choice in chess Amsterdam: Noord-Hollandsche Uitgeversmaatschappij.

DeValois RL and DeValois KK (1988) Spatial vision New York: Oxford University Press.

Diamond R, & Carey S (1986). Why faces are and are not special: An effect of expertise. J Exp Psychol Gen, 115(2), 107–117. [PubMed: 2940312]

Dosher BA, & Lu ZL (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. Proc Natl Acad Sci U S A, 95(23), 13988–13993. [PubMed: 9811913]

Duncker K (1945). On problem solving (translated by Lees LS), Psychological Monographs, 58, No. 270.

Egan DE, & Schwartz BJ (1979). Chunking in recall of symbolic drawings. Memory & Cognition, 7(2), 149–158. [PubMed: 88658]

Fahle M, & Edelman S (1993). Long-term learning in vernier acuity: Effects of stimulus orientation, range and of feedback. Vision Research, 33, 397–412. [PubMed: 8447110]

Fahle M, Edelman S, & Poggio T (1995). Fast perceptual learning in hyperacuity. Vision Research, 35(21), 3003–3013. [PubMed: 8533337]

Fahle M, & Poggio T, , (2002) Perceptual Learning (MIT Press, Cambridge, MA).

Feitelson D & Razel M (1984). Word Superiority and Word Shape Effects in Beginning Readers, International Journal of Behavioral Development, Vol. 7, No. 3, 359–370.

Fiser J, & Aslin RN (2002). Statistical learning of new visual feature combinations by infants. Proceedings of the National Academy of Sciences, 99, 15822–15826.

Gagne RM (1983). Some issues in the psychology of mathematics instruction. Journal for Research in Mathematics Education, 14 (1), 7–18.

Garrigan P, & Kellman PJ (2008). Perceptual learning depends on perceptual constancy. Proc Natl Acad Sci U S A, 105(6), 2248–2253. [PubMed: 18250303]

Ghose GM, Yang T, & Maunsell JH (2002). Physiological correlates of perceptual learning in monkey v1 and v2. J Neurophysiol, 87(4), 1867–1888. [PubMed: 11929908]

Gibson JJ (1950). The Perception of the Visual World Boston: Houghton Mifflin.

Gibson EJ (1969). Principles of perceptual learning and development New York: Prentice-Hall.

Gibson JJ (1979). The ecological approach to visual perception Boston: Houghton-Mifflin.

Gibson EJ, Owsley CJ, Walker A, & Megaw-Nyce J (1979). Development of the perception of invariants: Substance and shape. Perception, 8(6), 609–619. [PubMed: 530802]

Gibson JJ & Gibson EJ (1955). Perceptual learning: Differentiation or enrichment? Psychological Review, 62(1), 32–41. [PubMed: 14357525]

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

Gomez RL, & Gerken L (2000). Infant artificial language learning and language acquisition. Trends in Cognitive Sciences, 4(5), 178–186. [PubMed: 10782103]

Gregory RL (1972). Eye and brain: The psychology of seeing (2nd ed.). London: Weidenfeld and Nicolson.

Grossberg S (1997). Cortical dynamics of three-dimensional figure-ground perception of two-dimensional pictures. Psychological Review, 104, 618–58. [PubMed: 9243966]

Grossberg S (2003). How does the cerebral cortex work? Learning, attention, and grouping by the laminar circuits of visual cortex. Spatial Vision, 12, 163–85.

Guerlain S, La Follette M, Mersch TC, Mitchell BA, Poole GR, Calland JF, Jianhong Lv, & Chekan EG (2004). Improving surgical pattern recognition through repetitive viewing of video clips. IEEE Transactions on Systems, Man, and Cybernetic - Part A: Systems and Humans 34(6) 699–707.

Honger M-O, de Meneses YL, Beyeler A, & Jacot J, (2003). The resonant retina: Exploiting vibration noise to optimally detect edges in an image. IEEE Transactions on Pattern Analysis and Machine Intelligence, 25(9), 1051–1062.

Held R (1985). Binocular vision: Behavioral and neuronal development. In Mehler J & Fox R (Eds.), Neonate cognition: Beyond the blooming buzzing confusion (pp. 37–44). Hillsdale, NJ: Erlbaum.

Helmholtz H von (1864/1962). Treatise on physiological optics, Vol. III Edited by Southall JPC. NY: Dover Publications, Inc.

Hennig MH, Kerscher NJ, Funke K and Wo F (2002). Stochastic resonance in visual cortical neurons: Does the eye tremor actually improve visual acuity? Neurocomputing, 44, 115–120.

Hirsh-Pasek K & Golinkoff R (1997). The origins of grammar Cambridge, MA: The MIT Press.

Hubel DH, & Wiesel TN (1968). Receptive fields and functional architecture of monkey striate cortex. J Physiol, 195(1), 215–243. [PubMed: 4966457]

Huber L (1999) chapter in Fagot J (Ed.). Object recognition in animals, pp. 219–261.

Hull CL (1920). Quantitative aspects of the evolution of concepts. Psychological Monographs, XXVIII(1123), 1–86.

James W (1890). The principles of psychology, Vol I, New York: Dover Publications, Inc.

Johansson G (1973). Visual perception of biological motion and a model for its analysis. Perception & Psychophysics, 14, 201–211.

Jung CG (1923) Psychological Types (Baynes HB, Trans.). New York: Harcourt, Brace.

Karni A, & Sagi D (1991). Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. Proc Natl Acad Sci U S A, 88(11), 4966–4970. [PubMed: 2052578]

Karni A, & Sagi D (1993). The time course of learning a visual skill. Nature, 365(6443), 250–252. [PubMed: 8371779]

Keller FS (1958). The phantom plateau. Journal of the Experimental Analysis of Behavior, 1, 1–13. [PubMed: 14454938]

Kellman PJ (2002). Perceptual learning. In Stevens' Handbook of Experimental Psychology, Third Edition, Vol.3, Pashler H. H, and Gallistel CR, eds. New York: John Wiley & Sons, pp. 259–299.

Kellman PJ, & Arterberry ME (1998). The cradle of knowledge: Development of perception in infancy Cambridge, MA: The MIT Press.

Kellman PJ, Burke T, & Hummel J (1999). Modeling the discovery of abstract invariants. In Stankewicz B & Sanocki T (Eds.). Proceedings of the 7th Annual Workshop on Object Perception and Memory (OPAM), 48–51.

Kellman PJ, & Kaiser MK (1994). Extracting object motion during observer motion: Combining constraints from optic flow and binocular disparity. Journal of the Optical Society of America, A, Optics, Image Science & Vision, 12(3), 623–625.

Kellman PJ, Massey CM, Roth Z, Burke T, Zucker J, Saw A, Aguero K & Wise J (in press). Perceptual learning and the technology of expertise: Studies in fraction learning and algebra. Pragmatics and Cognition: Special Issue on Cognition and Technology

Kellman PJ & Spelke ES (1983). Perception of partly occluded objects in infancy. Cognitive Psychology, 15(4), 483–524. [PubMed: 6641127]

Kendler HH & Kendler TS (1961). Effects of verbalization on reversal shifts in children. Science, 134, 1619–1620. [PubMed: 14455126]

Koffka K (1935). Principles of gestalt psychology New York: Harcourt, Brace and World.

Kourtzi Z & Kanwisher N (2001) Representation of perceived object shape by the human lateral occipital complex. Science, 293(5534), 1506 – 1509. [PubMed: 11520991]

Kundel H & Nodine CF (1975). Interpreting chest radiographs without visual search. Radiology, 116, 527–532. [PubMed: 125436]

Leonards U, Rettenback R, Nase G, & Sireteanu R (2002). Perceptual learning of highly demanding visual search tasks. Vision Research, 42(18), 2193–2204. [PubMed: 12207979]

Lesgold A, Rubinson H, Feltovich P, Glaser R, & Klopfer D (1988). Expertise in a complex skill: Diagnosing x-ray pictures. In Chi MTH, Glaser R, & Farr M (Eds.), The nature of expertise Hilsdale, NJ: Lawrence Erlbaum Associates 311–342.

Liu Z (1999). Perceptual learning in motion discrimination that generalizes across motion directions. Proc Natl Acad Sci U S A, 96(24), 14085–14087. [PubMed: 10570202]

Liu T-T, Mayer-Kress G., and Newell KM (2006). Qualitative and quantitative change in the dynamics of motor learning. Journal of experimental psychology. Human perception and performance 32(2), 380–93. [PubMed: 16634677]

Locke J (1690/1971). An essay concerning human understanding New York: World Publishing Co.

Marcus GF, Vijayan S, Rao SB, & Vishton PM (1999). Rule learning by seven-month-old infants. Science, 283(5398), 77–80. [PubMed: 9872745]

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

Meltzoff AN & Moore MK (1977). Imitation of facial and manual gestures by human neonates. Science, 198, 75–78, [PubMed: 17741897]

Merzenich MM, Jenkins WM, Johnston P, Schreiner C, Miller SL, et al. (1996). Temporal processing deficits of language-learning impaired children ameliorated by training. Science, 271(5245), 77–81. [PubMed: 8539603]

Merzenich MM, Nelson RJ, Kaas JH, Stryker MP, Jenkins WM, Zook JM, Cynader MS, & Schoppmann A (1987). Variability in hand surface representations in areas 3b and 1 in adult owl and squirrel monkeys. Journal of Computational Neurology, 258, 281–286.

Michotte A (1952). Nuevos aspectos de la psicologia de la percepcion. Revista de Psicologia General y Aplicada, 7, 297–327.

Miller GA (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, 63, 81–97. [PubMed: 13310704]

Mollon JD & Danilova MV (1996). Three remarks on perceptual learning. Spat Vis, 10(1), 51–58. [PubMed: 8817771]

Nagarajan SS, Blake DT, Wright BA, Byl N, & Merzenich MM (1998). Practice-related improvements in somatosensory interval discrimination are temporally specific but generalize across skin location, hemisphere, and modality. Journal of Neuroscience, 18(4), 1559–1570. [PubMed: 9454861]

Neumann H, Yazdanbaksh A, & Mingolla E (2007). Seeing surfaces: The brain's vision of the world. Physics of Life Reviews, 4, 189–222.

Norcia AM & Tyler CW (1985). Spatial frequency sweep VEP: Visual acuity during the first year of life. Vision Research, 25(10), 1399–1408. [PubMed: 4090273]

Pasupathy A & Connor CE (2001). Shape representation in area V4: Position-specific tuning for boundary conformation. Journal of Neurophysiology, 86, 2505–2519. [PubMed: 11698538]

Perlovsky LI (2006). Toward Physics of the Mind: Concepts, Emotions, Consciousness, and Symbols. Physics of Life Reviews, 3(1), 22–55.

Petrov AA, Dosher BA, & Lu ZL (2005). The dynamics of perceptual learning: An incremental reweighting model. Psychol Rev, 112(4), 715–743. [PubMed: 16262466]

Piaget J (1952). The origins of intelligence in children New York: International Universities Press.

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

Pick AD (1965). Improvement of visual and tactual form discrimination. Journal of Experimental Psychology, 69(4), 331–339. [PubMed: 14286300]

Poggio T, Fahle M, & Edelman S (1992). Fast perceptual learning in visual hyperacuity. Science, 256(5059), 1018–1021. [PubMed: 1589770]

Purves D, Williams SM, Nundy S, & Lotto RB (2004). Perceiving the intensity of light. Psychological Review 111(1), 142–158. [PubMed: 14756591]

Recanzone GH, Merzenich MM, Jenkins WM, & Grajski KA (1992). Topographic reorganization of the hand representation in cortical area 3b of owl monkeys trained in a frequency-discrimination task. Journal of Neurophysiology, 67(5), 1031–1056. [PubMed: 1597696]

Recanzone GH, Schreiner CE, & Merzenich MM (1993). Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. Journal of Neuroscience, 13(1), 87–103. [PubMed: 8423485]

Reicher GM (1969). Perceptual recognition as a function of meaningfulness of stimulus material. Journal of Experimental Psychology, 81(2), 275–280.Saffran, J.R. (2002). [PubMed: 5811803]

Robinson CS, and Hayes JR (1978). Making inferences about relevance in understanding problems In Human Reasoning,Revlin R and Mayer RE, eds. Washington, DC: Winston.

Ryan TA & Schwartz CB (1956). Speed of perception as a function of mode of representations. Amer. J. Psychology, 69, 60–69.

Saarinen J & Levi DM (1995). Perceptual learning in vernier acuity: What is learned? Vision Research, 35(4), 519–527. [PubMed: 7900292]

Saffran JR, & Griepentrog GJ (2001). Absolute pitch in infant auditory learning: Evidence for developmental reorganization. Developmental Psychology, 37(1), 74–85. [PubMed: 11206435]

Saffran JR, Loman MM, & Robertson RRW (2000). Infant memory for musical experiences. Cognition, 77, 15–23.

Samuels SJ and Flor, Richard F (1997). The importance of automaticity for developing expertise in reading. Reading & Writing Quarterly, 13(2), 107 – 121.

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

Schoups A, Vogels R, Qian N and Orban G (2001). "Practising orientation identification improves orientation coding in V1 neurons. Nature, 412(6846), 549–553. [PubMed: 11484056]

Seitz A & Watanabe T (2003). Is subliminal learning really passive?, Nature, 22, 6927–6936.

Seitz A & Watanabe T (2005). A unified model of task-irrelevant and task-relevant perceptual. learning. Trends in Cognitive Science, Vol 9(7) 329–334.

Shiu LP, & Pashler H (1992). Improvement in line orientation discrimination is retinally local but dependent on cognitive set. Percept Psychophys, 52(5), 582–588. [PubMed: 1437491]

Silva AB & Kellman PJ (1999). Perceptual learning in mathematics: The algebra-geometry connection. In Hahn M & Stoness SC (Eds.). Proceedings of the Twenty-First Annual Conference of the Cognitive Science Society, Mahwah, NJ: Lawrence Erlbaum Associates, 683–688.

Simon HA (2001). Observations on the Sciences of Science Learning. Journal of Applied Developmental Psychology, 21, 115–121.

Sireteanu R & Rettenback R (2000). Perceptual learning in visual search generalizes over tasks, locations, and eyes. Vision Research, 40, 2925–2949. [PubMed: 11000393]

Slater A Mattock A & Brown E (1990). Size constancy at birth: Newborn infants' responses to retinal and real size. Journal of Experimental Child Psychology, 49(2), 314– 322. [PubMed: 2332727]

Spelke ES (2003). Developing knowledge of space: Core systems and new combinations. In Kosslyn SM & Galaburda A (Eds.), Languages of the Brain Cambridge, MA: Harvard Univ. Press.

Stickgold R, James L, and Hobson JA (2000). Visual discrimination learning requires sleep after training. Nature Neuroscience, 3, 1237–1238. [PubMed: 11100141]

Szokolszky A. (2003). An Interview With Eleanor Gibson, Ecological Psychology, 15(4), 271—281.

Tallal P, Merzenich M, Miller S, & Jenkins W (1998). Language learning impairment: Integrating research and remediation. Scandinavian Journal of Psychology, 39(3), 197–199. [PubMed: 9800537]

Titchener EB (1902). A textbook of psychology New York: Macmillan.

Trabasso TR & Bower GH Presolution dimensional shifts in concept identification. Journal of Mathematical Psychology, 1966, 3, 163–173.

Treisman A, & Gelade G, 1980 A feature integration theory of attention. Cognitive Psychology, 12, 97–136. [PubMed: 7351125]

Tsutsui KI, Sakata H, Naganuma T & Taira M (2002). Neural correlates for perception of 3D surface orientation from texture gradient. Science, 298, 409–412. [PubMed: 12376700]

Vogels R, & Orban GA (1985). The effect of practice on the oblique effect in line orientation judgments. Vision Res, 25(11), 1679–1687. [PubMed: 3832592]

von Hornbostel EM (1927). The unity of the senses. Psyche, 28, 83–89.

Wallach H (1948). Brightness constancy and the nature of achromatic colors. J. Exp. Psychol, 38, 310–324. [PubMed: 18865234]

Wang Q, Cavanagh P, & Green M (1994). Familiarity and pop-out in visual search. Perception & Psychophysics, 56(5), 495–500. [PubMed: 7991347]

Watanabe T, Náñez JE. & Sasaki Y. (2001). Perceptual learning without perception. Nature 413, 844–848. [PubMed: 11677607]

Wertheimer M (1938). Laws of organization in perceptual forms. In Ellis W (Ed.), A source book of Gestalt psychology (pp. 71–88). London: Routledge & Kregan Paul (Original work published 1923).

Westheimer G (1975). Visual acuity and hyperacuity. Investigative Ophthalmology, 14, 570–572. [PubMed: 1150397]

Westheimer G & McKee SP (1978). Stereoscopic acuity for moving retinal images. Journal of the Optical Society of America, 68(4), 450–455. [PubMed: 671135]

Wheeler DD (1970). Processes in the visual recognition of words. Dissertation Abstracts International,. 31(2-B), 940.

Wickens TD (2002). Elementary signal detection theory Oxford: Oxford University Press.

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

Xerri C, Merzenich MM, Jenkins W, & Santucci S (1999). Representational plasticity in cortical area 3b paralleling tactualmotor skill acquisition in adult monkeys. Cerebral Cortex, 9(3), 264–276. [PubMed: 10355907]

Yang T, Maunsell JHR (2004) The effect of perceptual learning on neuronal responses in monkey visual area V4. Journal of Neuroscience 24:1617–1626. [PubMed: 14973244]

Young KG, Li RW, Levi DM, & Klein SA (2005) Interocular specificity in perceptual learning of a position discrimination task. Invest Ophthalmol Vis Sci, 46: 5631.
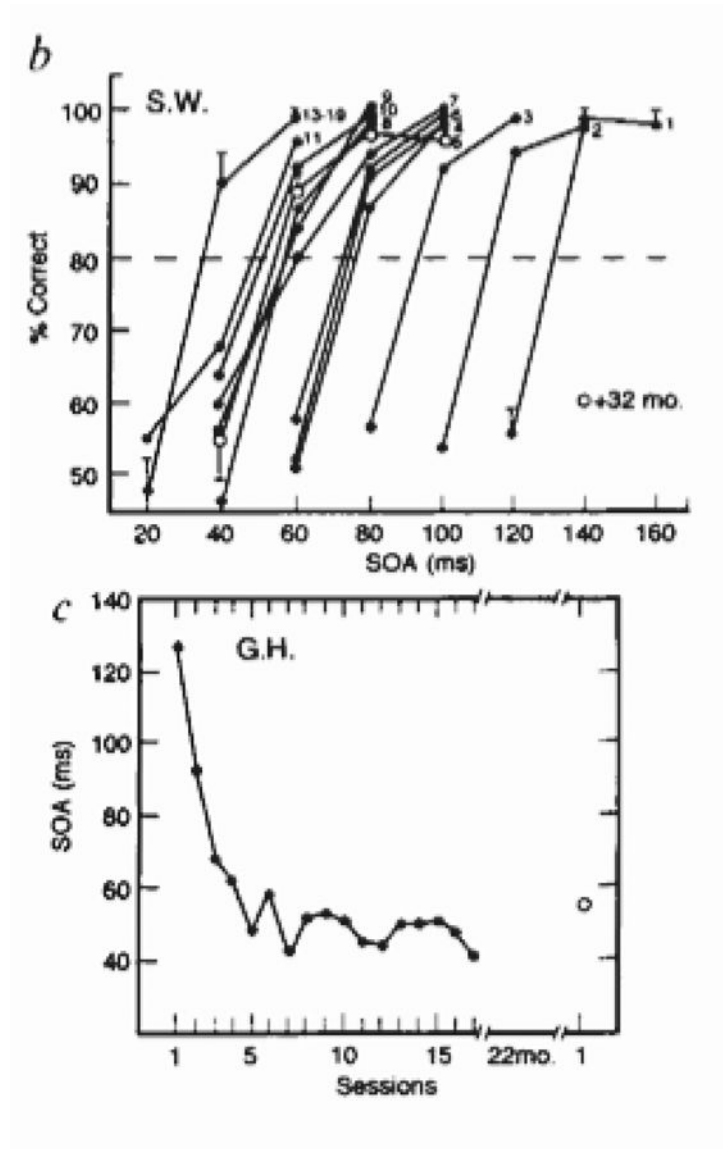
Figure 1.
Practice-induced improvements in visual search (Karni & Sagi, 1993). All data are from a single subject. Top: On each session (spaced 1–3 days apart), subjects showed significant improvement in the amount of exposure time (SOA) necessary to discriminate between two targets presented among distrators. Separate curves represent performance during different sessions. The leftward shift of the curves with practice indicates that performance of the task required progressively shorter presentations to achieve a given accuracy level. Between sessions, performance remained relatively constant, and was unchanged 32 months after the final training session. Bottom: Exposure time (SOA) necessary to reach threshold performance (~80% correct) as a function of training session.
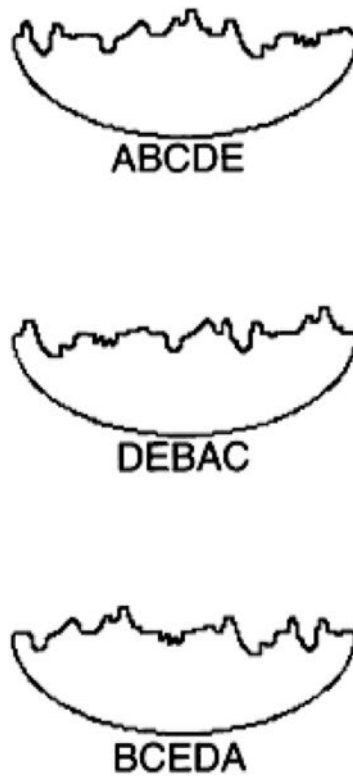
Figure 2.
Stimuli used in unitization experiment by Goldstone (2000). Each top contour is composed of 5 contour segments (labeled A-E). Categorization training in which a consistent arrangement (e.g., ABCDE) was discriminated from items differing by one component led to performance that indicated formation of a new perceptual unit or discovery of higher-order shape invariants (See text.)
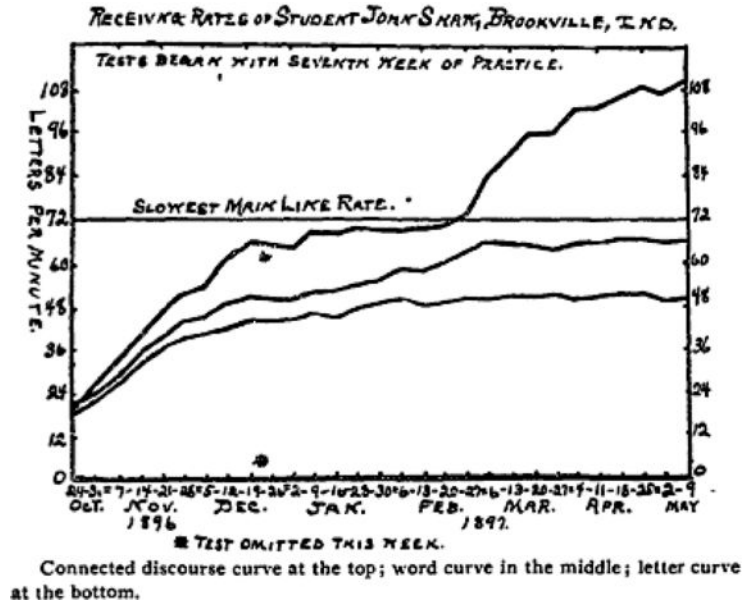
Fig XI.

RECEIVING RATES OF STUDENT JOHN SHAW, BROOKVILLE, IND.

TESTS BEGAN WITH SEVENTH WEEK OF PRACTICE.

SLOWEST MAIN LINE RATE.

Connected discourse curve at the top; word curve in the middle; letter curve at the bottom.

Figure 3.

Practice effects in telegraphic speech recognition (Morse Code). Subject performance increases with practice, then levels off for some time, with little apparent improvement with continued practice. During this time, letter recognition is becoming automatic. Automatic letter recognition frees up attentional resources for becoming sensitive to word structure, accounting for a second negatively accelerated learning curve (top line). Recognizing individual words without context (middle line) or individual letters that do not constitute words (bottom line), is always more difficult, and shows no indication of multiple learning stages.
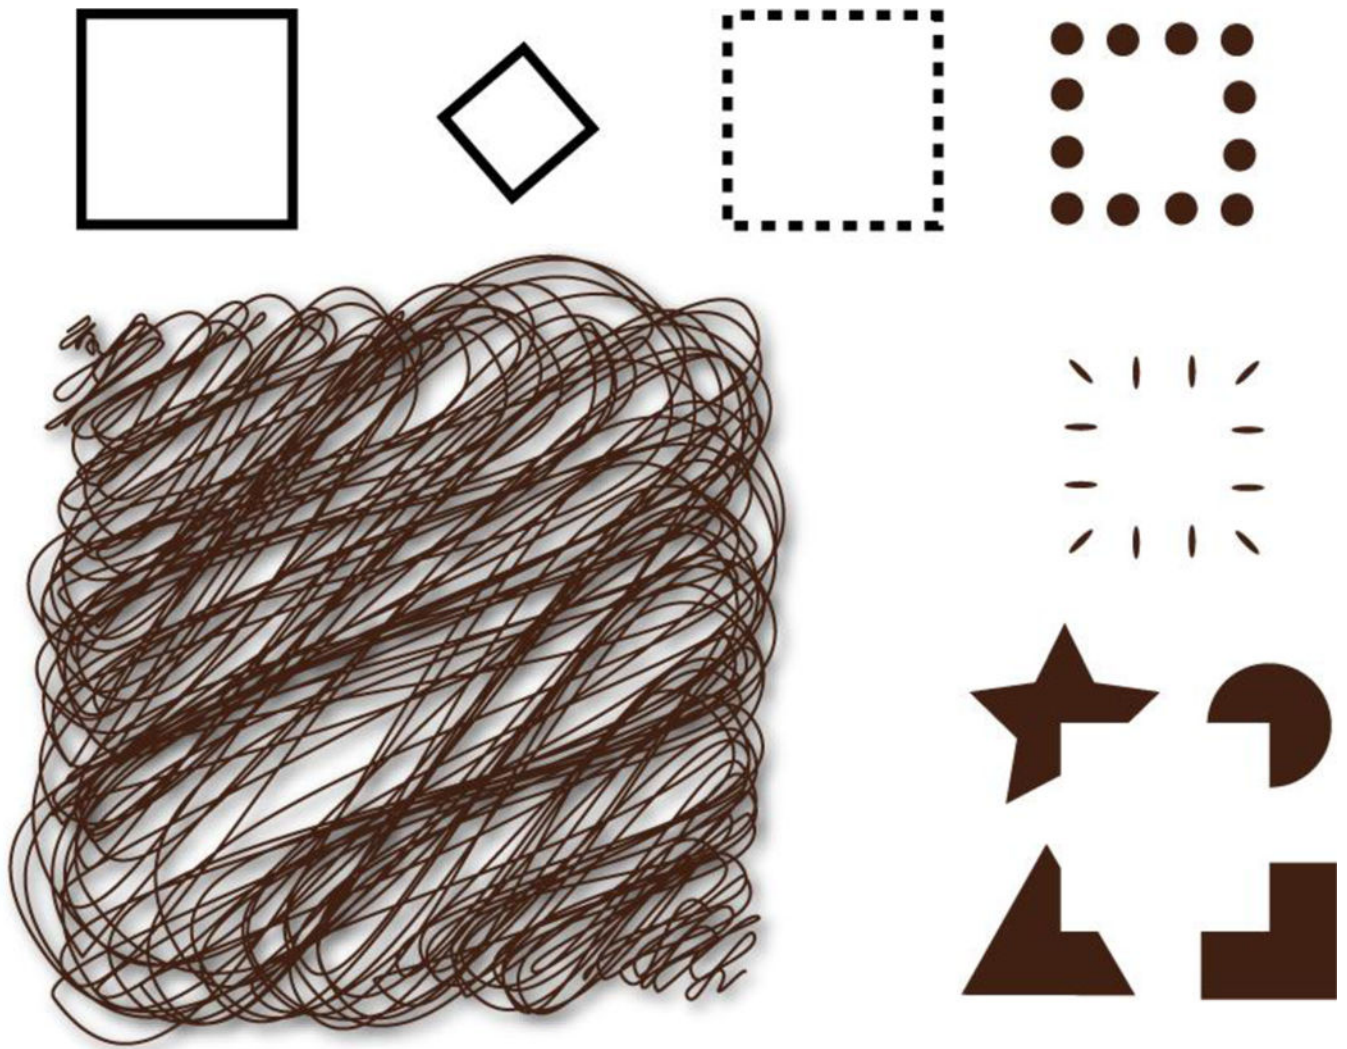
Figure 4.
Extraction of abstract invariants. All of these figures may be categorized as "squares", but none would activate a largely overlapping population of cells in visual cortex. What is the same about all of these figures is an abstract relation. Whatever is used to define a side, the sides are approximately equal. Specific concrete encodings from which this relation is extracted are incidental to squareness.
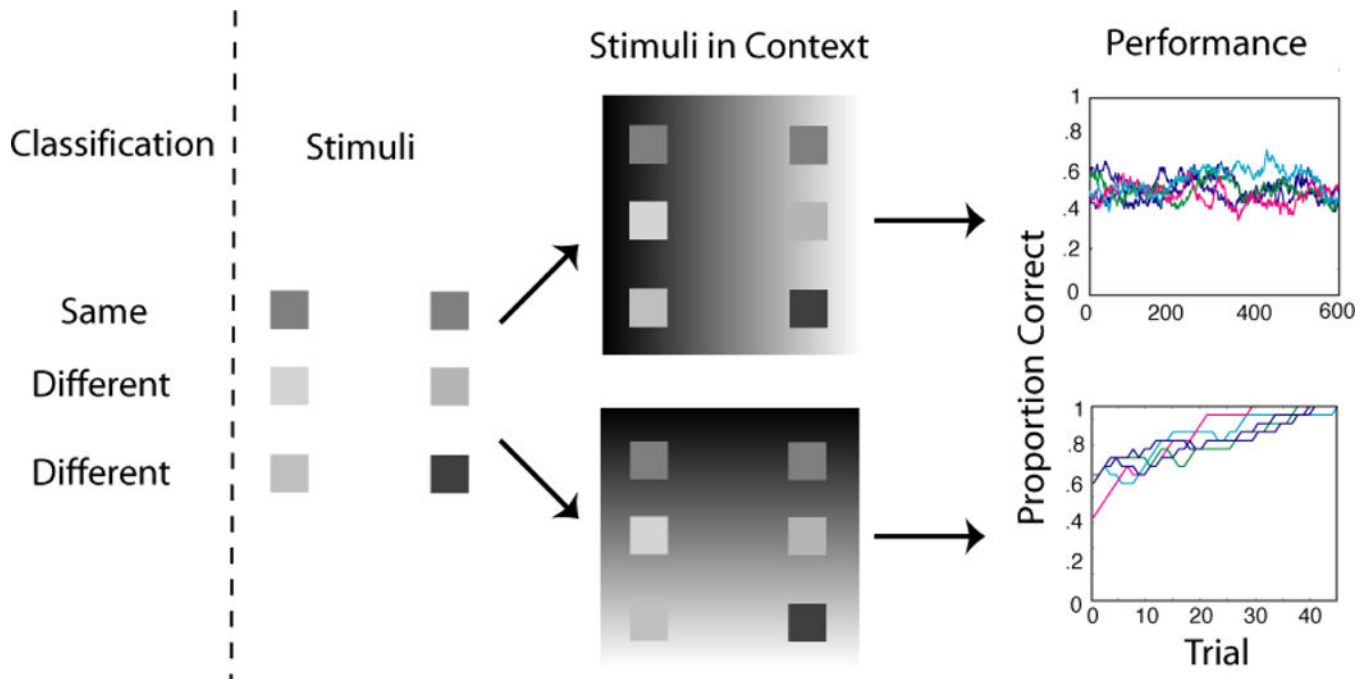
Figure 5.
Perceptual learning depends on perceptual constancy. Pairs of squares were shown on each trial, and the observer's task was to discover which ones were members of some category (see text). Perceived lightness depends on ratios of the local luminance (within a square) to its surround (Wallach, 1948). Bottom, center: When the invariant determining the classification was identical lightness of the two squares (accomplished by having the same local luminance for each and placing both on the same background), learning readily occurred. Acquisition curves are shown form 5 observers (bottom right). Top, center: When the identical local luminance squares were shown on a variable backgrounds, it caused differences in perceived lightness. In this condition, only the local luminances of the squares (a sensory invariant, decorrelated from perceived lightness) could be used to learn the classification. As in the other experiments involving size and motion, learning did not occur for any observer under this condition. These results suggest that perceptual learning works through constancy-based perceptual representations. (After Garrigan & Kellman, 2008).

Figure 6.

Table 1.

	NOVICE	EXPERT
DISCOVERY Effects		
SELECTIVITY:	Attention to irrelevant and relevant information	Selective pickup of relevant information/Filtering
UNITS:	Simple features	"Chunks"/ Higher-order relations
FLUENCY Effects		
SEARCH TYPE:	Serial processing	More parallel processing
ATTENTIONAL LOAD:	High	Low
SPEED:	Slow	Fast