

Perceptual learning depends on perceptual constancy

Patrick Garrigan* and Philip J. Kellman

Department of Psychology, St. Joseph's University, Philadelphia, PA19131; and Department of Psychology, University of California, Los Angeles, CA 90095-1563

Communicated by Charles R. Gallistel, Rutgers, The State University of New Jersey, Piscataway, NJ, December 17, 2007 (received for review January 29, 2007)

Perceptual learning refers to experience-induced improvements in the pick-up of information. Perceptual constancy describes the fact that, despite variable sensory input, perceptual representations typically correspond to stable properties of objects. Here, we show evidence of a strong link between perceptual learning and perceptual constancy: Perceptual learning depends on constancy-based perceptual representations. Perceptual learning may involve changes in early sensory analyzers, but such changes may in general be constrained by categorical distinctions among the high-level perceptual representations to which they contribute. Using established relations of perceptual constancy and sensory inputs, we tested the ability to discover regularities in tasks that dissociated perceptual and sensory invariants. We found that human subjects could learn to classify based on a perceptual invariant that depended on an underlying sensory invariant but could not learn the identical sensory invariant when it did not correlate with a perceptual invariant. These results suggest that constancy-based representations, known to be important for thought and action, also guide learning and plasticity.

abstract | representation

Classical theories and contemporary computational accounts of sensation and perception distinguish between variables encoded in early sensory analysis and higher level representations of objects, scenes, and events. Whereas early analyzers involve relatively local responses to energy, perceptual representations most often correspond to stable properties of material objects. Object properties persist across changes in the energy reaching the senses, so that comprehending the world requires perceptual constancy—attainment of relatively constant perceptual descriptions despite variation in the sensory inputs used to compute them. A common example is constancy of size: Under a variety of conditions, an object's perceived size does not vary as the observer's viewing distance changes, even though such changes alter the projected (retinal) size. Similarly, an object's surface lightness (shade of gray) does not appear to change when an object is viewed outside in sunshine or indoors, despite changes of more than three orders of magnitude in the light intensity reflected from that object to the eyes (lightness constancy). Perceptual constancies have often been claimed to involve learning, although evidence from human newborns has tended to disconfirm this idea (1). Here, we present evidence that perceptual constancy places strong constraints on learning.

Across modalities, tasks, and processing levels, perceptual learning (PL) plays a significant role in learning and expertise (2, 3). In recent years, however, PL research has focused largely on basic sensory discriminations; examples include Vernier acuity (4, 5), motion direction discrimination (6), and auditory frequency discrimination (7). This focus has helped illuminate connections between learning effects and neural plasticity, because the physiology of early sensory coding is both better understood and more accessible than that of higher-order representations. Improvement in simple sensory discriminations and physiological changes detected in early analyzers (8) have naturally led investigators to posit learning mechanisms directly based on early analyzers, such as those in primary sensory cortices (4). Likewise, findings that learning is specific to stimulus characteristics encoded early in processing, such as retinal

position (9) or orientation (10), have been interpreted as indicating early loci of learning.

It is also apparent, however, that learning involves interactions of lower processing levels with higher perceptual representations (11–13). Corresponding neurophysiological evidence suggests that higher cortical areas have important functional significance in PL and that the locus of neural modification for PL related to a single task depends on characteristics of the stimuli (14).

A framework for understanding such effects was proposed by Ahissar and Hochstein (15), who suggested that higher and lower levels are related via a “reverse hierarchy,” such that learning effects at higher levels precede and guide plasticity at lower levels. Specifically, learning takes place at the highest level at which the pertinent regularities exist. In many tasks, these regularities exist at relatively abstract levels, but when they do not, lower levels more directly related to the sensory input can be used for learning.

Here, we present evidence that these characteristics of PL involving higher and lower processing levels relate to a basic constraint: PL depends on perceptual constancy. Specifically, we tested the hypothesis that PL acts only through interpreted perceptual representations and cannot act directly on sensory inputs, even when task-relevant regularities are only present in those inputs.

We illustrate the approach, using the example of size perception. Under common conditions, perceived object size depends on a computation involving retinal (projective) size and distance information. From these inputs, the visual system computes a perceived size that corresponds well to real object size. Now consider an experiment in which PL is tested by having the learner discover the regularity that governs a classification (16). Research using this kind of task has been labeled both PL (2, 16) and category learning (17) in different research communities. We used this type of task both because the extraction of invariance from instances has been argued to involve the most ecologically important aspects of PL (2), and it allowed us to compare the accessibility of sensory and perceptual variables to learning processes.

Observers are shown displays and asked to respond “yes” or “no” as to whether the pair is a member of category X. The properties that determine the category are not described. The observers' task is to discover what information determines membership in the category based on accuracy feedback given after each response. In our size example, a pair of rectangles is presented on each trial, and the observer must answer “yes” or “no” as to whether the pair is a member of category X. On half of the trials, the two rectangles have identical size. The category is defined so that the correct answer is “yes” if and only if the two rectangles have the same size (Fig. 1 *Top*). Although the specific stimulus values change across trials, from outcome feedback,

Author contributions: P.G. and P.J.K. designed research; P.G. performed research; P.G. analyzed data; and P.G. and P.J.K. wrote the paper.

The authors declare no conflict of interest.

*To whom correspondence should be addressed. E-mail: patrick.garrigan@sju.edu.

This article contains supporting information online at www.pnas.org/cgi/content/full/0711878105/DC1.

© 2008 by The National Academy of Sciences of the USA




Fig. 1. For Experiments 1–3, sample stimuli for the correlated (categories defined by perceptual and sensory information) and uncorrelated (categories defined by sensory information alone) conditions are shown. Experiment (Exp.) 1, uncorrelated condition: Stimuli are stereopairs. “Same” rectangles had the same retinal size but were presented at different stereodisparities and therefore had uncorrelated perceived sizes. “Different” rectangles had different retinal sizes and uncorrelated perceived sizes. Experiment 1, correlated condition: Same rectangles had the same retinal sizes and correlated perceived sizes. Different rectangles had different retinal sizes and correlated perceived sizes. Experiment 2: This experiment had a similar design, with the shade of gray of the square determining category membership (again same or different) and perceived lightness either correlated or uncorrelated. Experiment 3: This experiment had a similar design, with retinal motion determining category membership (up or down) and perceived direction of motion either correlated or uncorrelated. Arrows indicate the possible directions of motion of each part of the stimulus. Experiment 4: This experiment had only one condition. Category membership was defined by perceptual, but not sensory equivalence. Two same response stimuli and one different response stimulus are shown. The same stimuli displayed are calibrated to the average perceived matching lightnesses across five subjects.

normal observers will gradually discover this regularity and classify accurately.

In this example, if the two rectangles in each pair are presented at the same observer-relative distance, this classification can be learned in either of two ways. The learner may discover that pairs in the category have equal perceived size or equal retinal size. (Because the members of the pair are equally distant from the observer, perceived and retinal sizes are correlated.) In a separate condition, however, we introduce stereoscopic depth differences between the two rectangles. Using this manipulation, pairs can have identical retinal sizes but different apparent distances. Differences in perceived distance change the rectangles’ perceived sizes but not their retinal sizes. The question is then: What regularities can be discovered in PL? Most studies of PL involve conditions of correlated sensory and perceptual information. Can an observer learn a category based on a sensory invariant but not a perceptual invariant (e.g., two rectangles that have the same retinal size but different perceived sizes)? We also asked whether an observer can learn a category based on a perceptual invariant without a sensory invariant.

We used this strategy—decoupling a sensory invariant from a perceptual invariant—in several perceptual domains (Fig. 1): perceived and retinal size (Experiment 1); perceived lightness and local brightness (Experiment 2); and perceived relative motion vs. absolute motion signals (Experiment 3).

In each domain, there are well established theories of how the sensory inputs we used are encoded and used to compute constancy-based perceptual representations. Perceived size is computed from an object’s retinal size and cues to its distance from the observer. When distance cues are removed, size estimates appear to be based to a large extent on the object’s visual angle. Perceived lightness in simple 2D scenes is derived from ratios of luminances of objects in a scene (18). In more complex scenes, other factors, such as the orientations of surfaces (19) and perceived illumination differences (20), can also affect perceived lightness. Direction-sensitive neural units are known to underlie motion perception (21), but perceived velocity of an object often depends on the relation of its local velocity to another object that acts as a reference. Using these relations, we tested the learnability of categories based on simple perceptual relations (involving perceived size, lightness, and motion) or simple relations of the sensory input from which these percepts are derived (retinal size, local brightness, and local motion).

These categories can be thought of as a subspace of a multidimensional feature-space, in which each classification is a point that lies within the subspace (in the category) or outside the subspace (not in the category). One measure of the ease of categorization is the simplicity with which one can place a boundary in the space that separates the sets of instances belonging to the two categories. When categories are linearly separable (i.e., the partitioning can be done with a straight line) categorization tends to be easy. We note here that the complexity of the learning task as defined by proximal stimulus dimensions or their corresponding sensory variables was matched. The difference in complexity between the two conditions exists only if the stimuli are encoded along perceptual, but not sensory, dimensions.

Results

Results were clear and opposite for relations defined by perceptual vs. sensory invariants (Fig. 2). For classifications defined by sensory invariants (e.g., equal retinal sizes but different perceived sizes), learning did not occur, even after hundreds of trials. These sensory invariants appeared to be undiscoverable by learning processes. In contrast, the identical sensory invariants were readily discoverable when they correlated with a perceptual classification (e.g., when equal retinal sizes correlated with equal perceived sizes). PL effects in learnable situations are normally largest in the earliest trials and conspicuously evident over the first few hundred trials (12, 13). One concern we had was that learning could have been occurring gradually in the cases where no reliable indications of performance improvement were observed. To address this possibility, we applied a sensitive algorithm for detecting changes in performance. This analysis allowed us to divide the data for each subject into two parts at the trial (the “change point”) for which the resulting subsets of the data were most consistent with two different levels of performance (22). Performance in the second subset of the data should therefore isolate those trials where subjects have learned or are beginning to learn. This analysis confirmed our initial conclusions. [See [supporting information \(SI\)](#).] Although there was strong evidence of a meaningful change point in all but one subject in the condition in which categories were defined by sensory and perceptual invariants (with performance averaging 85% in the latter subset), there was little evidence of meaningful change points among subjects in the condition in which learning could only be based on a sensory invariant (with no one attaining

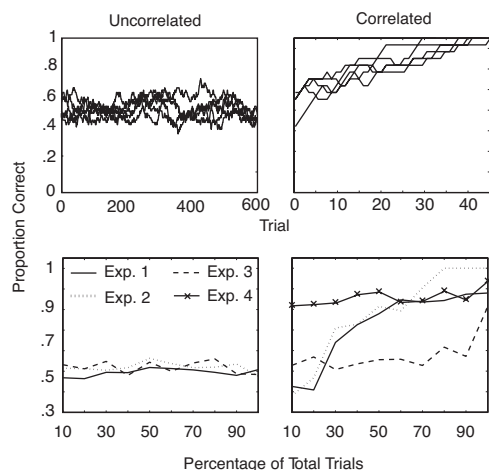


Fig. 2. Data are shown for individual subjects (*Upper*, Experiment 2) and averaged across subjects for each experiment (*Lower*). Results for categories defined by sensory information (Uncorrelated) are at *Left*. Results for categories defined by perceptual information (and sensory information; Experiments 1–3) are at *Right*. In the Uncorrelated condition, learning did not occur for any subject in any experiment. In the Correlated condition, learning is evident in all experiments. Learning was also evident in Experiment 4 (category defined by a perceptual, but not a sensory, invariant).

even 60% correct performance in the latter subset). Details of the procedure and statistical analysis are available in the [SI](#).

These results suggest that sensory variables are not directly accessible in learning; learning processes may use these inputs only through constancy-based representations. Although it remains possible that more extended training might lead to some learning, the lack of improvement in the sensory conditions across the entire session contrasts with many results showing that PL effects begin to appear early in training (12, 13). Moreover, our design used the same sensory relations in conditions that were unlearnable (in the uncorrelated conditions) as the basis of constancy in the correlated conditions. Learning was evident during early trials for all of the latter, indicating an important qualitative difference.

The conditions in which learning did occur contained both sensory and perceptual invariants; perhaps the combination facilitates learning more than either one alone. To assess this hypothesis, we conducted an additional experiment (Experiment 4), using local brightnesses and perceived lightnesses. In this experiment, the category to be learned was defined by perceptual (same shade of gray), but not sensory, invariance (different local brightnesses). Results showed that learning occurred for all subjects from perceptual invariance alone. Perceptual invariance derived from nonmatching sensory input was learnable, further supporting the notion that relations defined by constancy-based representations are key to learning, as opposed to some combination of perceptual and sensory invariance.

Discussion

These results suggest that PL may not be possible via direct access to sensory analyzers; it appears to be routed through perceptual classifications. In each of the perceptual domains tested, the sensory relations were unlearnable in a condition where they did not contribute to relevant perceptual classifications. These same sensory relations were readily learnable when they produced equivalent perceptual classifications. Remarkably, in experiments 1–3, the learnable perceptual classifications depended on sensory invariances that were by themselves unlearnable.

Perceptual vs. Compound Sensory Variables. One could argue that learning occurred based not on a truly perceptual variable but on a correlated compound sensory variable (in our perceived size experiment, for example, a certain set of values in a space defined by dimensions of retinal size and binocular disparity). Discrimination based on sets of values on multidimensional sensory variables could be hypothesized to give the same results as a perceptual variable (because perceptual variables are derived from relations in sensory inputs). There are reasons for considering such an alternative unlikely, however.

First, there is no independent evidence for sensory coding of the novel conjunctive sensory variables that would be required (such as coding particular retinal size–disparity combinations). Also, in one of our studies, identical perceptual values were achieved from many separate sensory combinations. It is straightforward to understand how learning occurred based on extraction of perceptual values, but understanding the same learning through conjunctively coded sensory variables is problematic, because such learning would seem to be too specific. It is not obvious how learning could generalize from one or several particular retinal size–binocular disparity pairings to others, except insofar as they signify the same perceived size.

Another issue is that the compound sensory variables required would need to be quite complicated. In the size–distance example, although retinal size and binocular disparity are both quantities that are likely computed in early visual processing, these alone would not suffice. An accurate surrogate for perceived size would require relative disparities scaled by the observer–relative distance, obtained from some other source, to at least one point in the scene (23). (Binocular disparities alone do not provide distance information.) In short, a compound sensory variable that could explain our results would essentially mimic the same complex computations that lead to perceived size.

In the domain of space perception, an additional insight about such computations is suggested by neurophysiological data. Computation of depth and slant appear to involve later visual areas, including parietal areas, such as cIPS (24), and temporal areas (25). A consistent finding about spatial processing in these areas is that neural responses can often be elicited by different cues to the same perceptual property (24, 25). Such findings fit more readily with computation of perceptual variables than with particular multidimensional sensory ones. There would be no reason to expect, for example, that a retinal size–disparity combination should be interchangeable with a retinal size–texture gradient combination, except insofar as these produce equivalent perceptual quantities.

Finally, although the results in our “perceptual” conditions could be claimed to involve in each case a novel and complicated sensory variable, this would make the pattern of results paradoxical. Recall that learning did not occur in our sensory conditions, where in each case a simple sensory variable known to be encoded in early processing governed the classification to be learned. If learning processes have access to sensory variables, it would be odd if that access were limited to novel complicated ones but excluded known simple ones. One could imagine that only complicated sensory variables that correspond to simple perceptual variables are learnable, but that idea would closely resemble (and be experimentally indistinguishable from) our proposal that perceptual variables guide learning.

More generally, the idea that perception cannot represent stimulus parameters encoded in early sensory responses is common to a variety of otherwise diverse theoretical views. There is a deep reason for this, based on the relation of matter and energy in perception. We perceive by means of energy received at the senses, but it is the properties of the material world—objects, surfaces, spatial arrangements, and events—that matter most for thought and action. Early responses in each sensory system necessarily relate to energy dimensions, but

obtaining perceptual attributes that reflect properties of the material world requires computing relations among sensory activations. The geometry of a surface, for example, may be important, but it can be derived in vision from numerous sources, such as contours, motion perspective, binocular disparity, shading and texture, each of which is itself a relational variable. The same surface geometry may also be perceived through other senses, such as kinesthesia or touch. Perceptual experience is largely concerned with the outputs of computations—outputs that represent environmental properties. What our results suggest is that, like perceptual experience, learning processes are constrained to use the outputs of perceptual computations.

That both perceptual experience and learning appear to be constrained to use the outputs of perceptual computations reflects both the ecological importance of perceptual regularities and issues of economy and efficiency. One could imagine a system in which perception and learning both have access to preliminary encodings that are used in perceptual computations; in fact, most models of PL have assumed such access. It is likely that such access would overload any information processing system. Sensory data fluctuate continually. The perceived surface color of a book one is carrying does not change as you walk with it, but the amount of light reflected to the eyes from any point on the book changes continually as the observer passes in and out of shadows, as the sun goes behind a cloud, or as the book's orientation changes even slightly in one's hand. Beside the information load of experiencing or learning about sensory fluctuations, there is also the issue of efficiency, in that most of these fluctuations do not encompass behaviorally relevant regularities. Even if they could be stored and accessed, they might make apprehension of important properties more difficult.

Percepts and Natural Scene Statistics. The point regarding the need for computations that extract relations can be made in the context of most theories of perception, but it has been sharpened by research suggesting important relations between perceptual outcomes and the statistics of natural scenes.

The correspondence between perception and scene statistics may indicate that perceptual systems have incorporated important constraints on the physical world or that actual probability distributions about the physical characteristics of the world and their appearances in different contexts are somehow encoded. Both of these options pose a problem of complexity in understanding how such regularities are acquired. For example, Yang and Purves (26) showed that a number of illusions of lightness perception that have been difficult to understand in terms of current theories could be predicted by sufficiently detailed statistics of image-source relations—capturing the range of brightnesses that are exhibited by particular surfaces in different configurations. In this domain and others, it has been argued, percepts depend on statistics about the mapping between proximal and distal stimuli across different contexts, likely incorporated over evolutionary time into perceptual computations (27). Yet they note that estimating these probability distributions remains a serious obstacle for statistical approaches to perceptual inference, because real world scenes are typically very complicated.

Yang and Purves (2003) (28) suggested an alternative approach to generating percepts in which sensory properties in similar contexts are ranked relative to one another. It is this dimensionally reduced ranking, not the full joint distribution, that determines perceptual appearance. Moreover, scene characteristics may steer perceptual processing, even from quite early levels toward the appropriate context needed to obtain the appropriate ranges of perceptual values (26).

We have used recent statistical views of perception as an illustration, but the key idea is likely to be a general feature of perceptual processing and perceptual theories. Sensory values

may be encoded at early stages, but relational processing derives higher-order regularities, and only the latter comprise perceptual representations and accessible inputs for learning. Early sensory encodings do not by themselves indicate the physical state of the world. Sensory encoding is indispensable for perception, but because of issues of economy and efficiency, or perhaps for other reasons, there appears to be little access to early sensory encodings in perception and PL.

Experimental Tasks and the Scope of PL. The experimental paradigm used here allowed us to test whether PL processes could access particular sensory variables when they did or did not lead to regularities in constancy-based perceptual representations. The task tested discovery of the stimulus information underlying a classification, which has been argued to be the most crucial aspect of PL (2).

Recently, it has been common for investigators to study PL, using simple discriminations, often restricted to two fixed stimuli, along with explicit instructions to subjects about the discrimination to be made. Confinement of the task in these ways has been claimed to allow inferences about the loci of learning effects, especially in connection with animal models showing plasticity in primary sensory cortices or with findings showing specificity of learning, i.e., lack of transfer across retinal locations or changes in stimulus attributes (e.g., orientation, motion direction, etc.) It has been explicitly suggested by some and tacitly accepted by many that “perceptual learning” should in fact be defined as involving only low level modifications in sensory systems, perhaps including only primary sensory cortices. For example, Fahle and Poggio (29) defined PL 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.” In vision, discrimination tasks with two fixed stimuli shown in a restricted retinal area naturally fit with this emphasis, because one might thereby address a restricted pool of units in the first cortical areas (V1, V2) known to be selective for specific retinal locations and stimulus attributes.

Such attempts to confine PL to a particular task or to primary sensory cortices are problematic, however. The notion that specificity of transfer implies a low-level locus of processing has been criticized as a fallacy (30, 31). Mollon and Danilova (30) argued that learning effects interpreted as changes in low level units are consistent with more central learning processes that discern which outputs from earlier levels are relevant to the task. This notion of selection not only accords with much earlier work on PL but characterizes recent models of low-level PL (32). It also coheres with the fact that actual data about transfer from PL experiments have been inconsistent (33, 34). Small task variations can lead to big differences in generality of transfer. Such differences [e.g., discriminating motion directions differing by 3° vs. 8° (34)] suggest that, whereas specificity of transfer is an interesting issue in PL, it should not be used to define PL. Given that specificity of transfer need not imply changes at a low level, there are few if any studies of humans that furnish any evidence for confinement of perceptual learning to sensory cortices.

A separate issue is that learning effects, if confined to primary sensory cortices, would likely not be perceptual learning. In vision, it is reasonably clear that processors in V1 and V2 are unable to furnish 3D spatial position (35), figure/ground assignment (36) or shape (37), and other perceptual properties that are probably computed further along.

PL effects important in real-world tasks often involve these properties (38). There may be a common recent tendency to equate perceptual learning with “sensory plasticity,” but these are not necessarily synonymous, especially if the consequence of equating them is to limit PL to the simplest sensory discriminations.

In our view, the kind of task used here has greater ecological

relevance than simple discrimination tasks. As humans encounter a range of individual examples of objects, situations, and events, they must discover the underlying regularities in the input that govern important behavioral consequences. Not only is this version of PL most relevant to the child's task in learning what a dog, square, or toy is, it is also known to be a crucial basis of advanced human expertise (3).

Perhaps the best reason, however, that PL research should not be confined to a single type of task is that artificial distinctions relating to paradigms may obscure underlying invariance in PL processes and mechanisms. The emphasis on early sensory cortices in PL is to some degree tied up with a particular idea about mechanism: PL effects might be changes in the receptive fields of cortical units or the number of units attuned to particular stimulus attributes at this level (6). In vision, however, PL tasks in monkeys have produced large behavioral improvements, but single-cell recording before and after has revealed little evidence of change in receptive field properties or recruitment of additional units at the earliest cortical levels (V1, V2). One study of orientation discrimination found some increases in the slopes of tuning curves of V1 units $\approx 20^\circ$ away from the trained orientation, but a subsequent study found no such effect nor any reliable effect of training, compared with control neurons, in any of eight receptive field parameters studied in V1 and V2 (39). Modest evidence of receptive field changes has been reported after training in visual area V4, but the changes are unlikely to be large enough to support the behavioral improvements observed.

The lack of clear evidence of changes found in the earliest visual cortical areas is consistent with the theoretical idea that PL changes in visual tasks primarily involve, not the modification of early analyzers, but selection by higher-level processes of the analyzer outputs from earlier levels that best determine the classification being trained. Recently, Petrov *et al.* (32) carried out experimental and modeling work to compare these two potential mechanisms of PL: the "representation modification" idea (e.g., changes in receptive fields of early analyzers) vs. a "selective reweighting" idea, in which PL consists of gradual learning of which analyzers are most useful for a given task.

Petrov *et al.* found their data in an orientation discrimination task could be completely accounted for by a model using only selective reweighting. Their detailed model was conspicuous in three respects: It was a fully functioning learning model that performed trial-by-trial learning with gray-scale images as inputs; it used well documented features of orientation-sensitive units to arrive at input representations and coupled these with a simple connectionist reweighting scheme; and, the model fit the data remarkably well, with no free parameters. These investigators suggest that their model of selection and reweighting of analyzers is consistent with most existing PL data. It is interesting to reflect that Gibson (2) argued that the essence of PL is selection and that what is learned in PL are distinguishing features, those aspects of the stimulus array that make the difference for a classification. Although Gibson's view was cast in terms of selection among stimulus features and the view of Petrov *et al.* in terms of selection among analyzers (along with a much more detailed model), they are in essence the same idea. (Because information must be encoded to be used, and the function of encoding processes is to obtain information, selection among analyzers and selection of stimulus information are notions not easily separated.) If so, there may be continuity between the simple discriminations used in some studies and the wider variety of PL tasks and stimulus contexts used in other PL work.

If this analysis is correct, the specificity of PL in different contexts will be a consequence of the experimenter-chosen task. Selective reweighting may occur at various levels, depending on where the regularities relevant to the task reside (15). This idea

connects results across tasks and stimulus domains, such as the elementary discrimination tasks often used in recent research and tasks using arguably more ecologically natural tasks and relations, such as in the work of Gibson (2) and in the work reported here. Our experimental results, in fact, both reflect this overall view of PL and indicate a significant constraint on it. In our experimental tasks across several perceptual domains, the selection and weighting of simple stimulus parameters could have led to accurate responding. But a crucial constraint on selection and weighting in any task will be what inputs are accessible. Our results suggest that stimulus relations given by sensory variables alone were not accessible to selection processes. When these same variables led to meaningful regularities in constancy-based representations, however, they were accessible, and learning based on selective extraction of these regularities readily occurred.

General Implications. These results place basic constraints on computational and physiological models of PL and make new predictions about what is learnable via PL. Learning accesses representations of task-relevant environmental properties. This makes sense ecologically: Most relevant for learning are regularities in the world rather than fluctuations of energy at sensory surfaces. In real-world tasks, the important regularities for learning may seldom be explicit in early representations that do not incorporate perceptual constancy. Constraining PL in this way adaptively limits the number of relationships that must be considered as potentially important regularities requiring further processing by the brain.

Our hypothesis does not deny the importance for learning of information at early sensory processing stages or even that learning may include physiological changes at early levels; rather, it implies that the use of early information and low-level changes are driven by perceptual classifications. Thus, our view is consistent with that of Ahissar and Hochstein (15), who proposed a "reverse hierarchy" theory of PL: the idea that "learning is a top-down process, which begins at high-level areas of the visual system, and when these do not suffice, progresses backwards to the input levels. . ." (15). We add to this view that preconstancy sensory regularities may be inaccessible, even when they are present in the system and provide the only means of performing a task.

The idea that learning must be guided by constancy-based representations implies that particular stimulus regularities can be used in PL if they are used to achieve perceptual classifications, but these very same regularities cannot be used otherwise. This idea about constraints on perceptual learning would seem to apply quite generally to learning. Our paradigm, for example, involves discovery of relations important to a classification but also connecting them to a response or label. The latter associative component may be relatively trivial in these examples, but it raises an important point. Like PL, associative learning may also be constrained by perceptual constancy. To be a "stimulus," it may not be sufficient that there exists some sensory registration within the organism. Rather, certain outputs of perceptual processes—constancy-based representations—likely comprise the domain of available stimuli, whether in the perceptual discovery of important relations or in learning by relating environmental attributes to each other and to behavior.

The current results bear interesting relations to recent findings that PL can occur without awareness. At first blush, it may seem that the dependence of learning on constancy-based representations is inconsistent with learning based on subthreshold stimuli [e.g., subthreshold motion signals (40, 41)]. Our experiments distinguish the sensory and perceptual, whereas these experiments distinguish subthreshold and suprathreshold. It is an open and interesting question whether subthreshold signals are

processed into representations of properties of the world, just like suprathreshold signals.

Our results and the theorized relation between constancy and learning may also have implications for the classical debate about learning of perceptual constancies. Although our work does not test these issues developmentally, it raises the possibility that the classical account—that constancy emerges from associating sensory experiences with each other and with action—is impossible. If the present findings about learnability apply early in development, relations among purely sensory inputs that do not work through higher-order perceptual representations would be unlearnable. This perspective would be consistent both with evidence indicating meaningful perception from birth in humans and also with the view that learning in any perceptual domain builds by correlation with representations furnished by at least one unlearned perceptual process, a position proposed originally by Wallach (42). Although there are many opportunities for learning in perception, discovery of relations from uninterpreted sensory inputs may not be among them.

Methods

General Methods. All experiments were carried out in a dark room, using images presented on a Viewsonic CRT monitor with an 1152 × 870 pixel

resolution. Subjects were positioned in a headrest 40 cm from the center of the monitor screen. Audio feedback indicating correct and incorrect responses was given for all experiments. Before each experiment, all subjects were given a practice task in which they were instructed to learn a category. On each trial of the practice session, two shapes were shown to the subjects. Each stimulus was in the category (correct response: “yes”) if both shapes had the same color, and was not in the category (correct response: “no”) if the shapes had different colors. Practice continued until subjects responded with the correct categorization on 90% of their last 20 trials. In all experiments, subjects were instructed to try to learn the category, and that if they made a certain number of consecutive correct responses (the learning criterion, known to all subjects), the experiment would end. In a randomly chosen half of all scheduled trials, the stimulus presented was a member of the category, and, in the remaining trials, the stimulus was not a member of the category. Five subjects participated in each condition of each experiment. The maximum number of trials for Experiments 1–4 was 756, 800, 500, and 800, respectively.

ACKNOWLEDGMENTS. We thank B. Backus, C. R. Gallistel, C. Massey, J. Nanez, D. Purves, A. Seitz, and several anonymous reviewers for insightful comments. This work was supported by National Science Foundation Grant ROLE-0231826 and National Eye Institute Grant EY13518 (to P.J.K.) and National Science Foundation Grant IBN-0344678.

- Slater A, Mattock A, Brown E J (1990) Size constancy at birth: Newborn infants' responses to retinal and real size. *J Exp Child Psychol* 49:314–322.
- Gibson E (1969) *Principles of Perceptual Learning and Development* (Appleton, New York).
- Kellman PJ (2002) Perceptual learning. In *Stevens' Handbook of Experimental Psychology*, eds Pashler H, Gallistel CR (John Wiley & Sons, New York), 3rd Ed, Vol 3, pp 259–299.
- Poggio T, Fahle M, Edelman S (1992) Fast perceptual learning in visual hyperacuity. *Science* 256:1018–1021.
- McKee SP, Westheimer G (1978) Improvement in Vernier acuity with practice. *Percept Psychophys* 24:258–262.
- Ball K, Sekuler R (1982) A specific and enduring improvement in visual motion discrimination. *Science* 218:697–698.
- Recanzone GH, Schreiner CE, Merzenich MM (1993) Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *J Neurosci* 13:87–103.
- Wang X, Merzenich MM, Sameshima K, Jenkins WM (1995) Remodeling of hand representation in adult cortex determined by timing of tactile stimulation. *Nature* 378:71–75.
- Fahle M (2004) Perceptual learning: A case for early selection. *J Vision* 4:879–890.
- Shoups AA, Vogels R, Orban GA (1995) Human perceptual learning in identifying the oblique orientation: Retinotopy, orientation specificity and monocularly. *J Physiol* 483:797–810.
- Ahissar M, Hochstein S (1993) Attentional control of early perceptual learning. *Proc Natl Acad Sci USA* 90:5718–5722.
- Karni A, Sagi D (1993) The time course of learning a visual skill. *Nature* 365:250–252.
- Haijiang Q, Saunders JA, Stone RW, Backus BT (2006) Demonstration of cue recruitment: Change in visual appearance by means of Pavlovian conditioning. *Proc Natl Acad Sci USA* 103:483–488.
- Song Y, Ding Y, Fan S, Qu Z, Xu L, Lu C, Peng D (2005) Neural substrates of visual perceptual learning of simple and complex stimuli. *Clin Neurophysiol* 116:632–639.
- Ahissar M, Hochstein S (1997) Task difficulty and the specificity of perceptual learning. *Nature* 387:401–406.
- Goldstone RL (2000) Unitization during category learning. *J Exp Psychol Hum Percept Perform* 26:86–112.
- Ashby FG, Maddox WG (2005) Human category learning. *Annu Rev Psychol* 56:149–178.
- Wallach H (1948) Brightness constancy and the nature of achromatic colors. *J Exp Psychol* 38:310–324.
- Gilchrist A (1980) When does perceived lightness depend on perceived spatial arrangement? *Percept Psychophys* 28:527–538.
- Howe PDL (2006) Testing the coplanar ratio hypothesis of lightness perception. *Perception* 35:291–301.
- Newsome WT, Shadlen MN, Zohary E, Britten KH, Movshon JA (1995) *The Cognitive Neurosciences*, ed Gazzaniga MS (MIT Press, Cambridge, MA) pp 401–414.
- Gallistel CR, Balsam PD, Fairhurst S (2004) The learning curve: Implications of a quantitative analysis. *Proc Natl Acad Sci USA* 101:13124–13131.
- Wallach H, Zuckerman C (1963). The constancy of stereoscopic depth. *Am J Psychol* 76:404–412.
- Tsutsui KI, Sakata H, Naganuma T, Taira M (2002) Neural correlates for perception of 3D surface orientation from texture gradient. *Science* 298:409–412.
- Liu Y, Vogels R, Orban GA (2004) Convergence of depth from texture and depth from disparity in macaque inferior temporal cortex. *J Neurosci* 24(15):3795–3800.
- Yang Z, Purves D (2004) The statistical structure of natural light patterns determines perceived light intensity. *Proc Natl Acad Sci USA* 101:8745–8750.
- Howe CQ, Lotto RB, Purves D (2006) Empirical approaches to understanding visual perception. *J Theor Biol* 241:866–875.
- Yang Z, Purves D (2003) Image/source statistics of surfaces in natural scenes. *Network Comput Neural Sys* 14:371–390.
- Fahle M, Poggio T, eds (2002) *Perceptual Learning* (MIT Press, Cambridge, MA).
- Mollon JD, Danilova MV (1996) Three remarks on perceptual learning. *Spatial Vision* 10:51–58.
- Dosher B, Lu Z-L (1998) Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proc Natl Acad Sci USA* 95:13988–13993.
- Petrov A, Dosher B, Lu Z-L (2005) The dynamics of perceptual learning: An incremental reweighting model. *Psychol Rev* 112(4):715–743.
- Fahle M, Edelman S, Poggio T (1995) Fast perceptual learning in hyperacuity. *Vision Res* 35:3003–3013.
- Liu Z (1999) Perceptual learning in motion discrimination that generalizes across motion directions. *Proc Natl Acad Sci USA* 96(24):14085–14087.
- Cumming BG, Parker AJ (1997) Responses of primary visual cortical neurons to binocular disparity without depth perception. *Nature* 389:280–282.
- Baylis G, Driver J (2001) Nat Neurosci 4, 937–942 (2001) Shape-coding in IT cells generalizes over contrast and mirror reversal, but not figure-ground reversal. *Nat Neurosci* 4:937–942.
- Kourtzi Z, Kanwisher N (2001) Representation of perceived object shape by the human lateral occipital complex. *Science* 293(5534):1506–1509.
- Diamond R, Carey S (1986) Why faces are not special: an effect of expertise. *J Exp Psychol* 115:107–117.
- Ghose GM, Yang T, Maunsell JHR (2002) Physiological correlates of perceptual learning in monkey V1 and V2. *J Neurophysiol* 87:1867–1888.
- Watanabe T, Sasaki Y, Nanez J (2001) Perceptual learning without perception. *Nature* 413:844–848.
- Seitz AR, Watanabe T (2003) Is subliminal learning really passive? *Nature* 422:36.
- Wallach H (1985) Learned stimulation in space and motion perception. *Am Psychol* 40:399–404.