

Oxford Handbooks Online

Perceptual Learning

Robert L. Goldstone and Lisa A. Byrge

The Oxford Handbook of Philosophy of Perception (*Forthcoming*)

Edited by Mohan Matthen

Online Publication Date: Sep
2013

Subject: Philosophy, Philosophy of Mind, Epistemology

DOI: 10.1093/oxfordhb/9780199600472.013.029

[–] Abstract and Keywords

This article has been commissioned as part of the forthcoming *Oxford Handbook of Philosophy of Perception* edited by Mohan Matthen. Through perceptual learning, perceptual systems are gradually modified so as to better fit an organism's environment and frequently occurring needs. We consider psychological and neurophysiological evidence that changes to perception can be early in the stream of information processing. Three specific mechanisms of perceptual learning are described: attentional tuning, unitization, and attribute differentiation. These mechanisms allow organisms to emphasize important perceptual information, to construct single functional units that are activated when a familiar complex configuration arises, and to isolate perceptual attributes that were originally psychologically fused. We describe ways by which people modify their perceptual systems so as to better meet their goals, and the implications of these modifications for the cognitive penetrability of perception, relations between perception and higher-order reasoning, and education.

Keywords: perceptual learning, cognitive penetrability, neural plasticity, categorical perception, unitization, attention, education, modularity

Perception can be learned. Experience shapes the way people see and hear. In one sense, these are barely interesting claims. After all, experience provides the sensory input to our perceptions as well as knowledge about the identities and functions of the objects that make up our physical environment. Perceptual learning, however, speaks to a much deeper relationship between experience and perception, in which fundamentally different perceptions of the same sensory input may arise in individuals with differing experiences or training. This raises important issues about the ontology of sensory experience, the relationship between cognition and perception, and the possibility of a theory-neutral perceptual ground for science.

Given the importance of responding efficiently and effectively to the environment, one might expect that 'hard-wired' perceptual circuitry would be an optimal design, especially given that the optic properties of the environment are largely stable. Plastic perceptual circuitry might be thought to be risky, given that changes made to the early representation of information will affect all subsequent processing. 'Early' in the previous sentence has two meanings, and both are relevant to our discussion. Developmentally early perceptual changes are those that occur during infancy. Operationally early perceptual changes are those that occur during the first few processing stages of transducing external signals to brain events.

Although developmentally and operationally early perceptual changes are risky, ubiquitous variation in every individual's environment—in local and regional attributes of objects, flora, fauna, even communication systems—highlight the inherent limitations of 'built-in' perceptual mechanisms and the importance of mechanisms for perceptual learning. One might expect that developmentally early perceptual learning would occur rarely, because evolution should have already tuned our perceptual systems to be sensitive to the most important elements of our shared environment. In fact, developmental research shows that some human perceptual constraints appear to be learned rather than innate (Quinn and Bhatt, 2006). This research indicates that the observed perceptual plasticity is often times highly specific to the trained environment rather than the environment acting as a general trigger for

the maturation of the perceptual system. For example, 4-month old infants develop perceptual representations for specific configurations of visual elements that co-occur (Needham et al., 2005).

Furthermore, work in animals (Blakemore and Cooper, 1970) has shown that the absence of environmental regularities (such as horizontally orientated features) early in development can result in the lack of neurons dedicated to these features in the primary visual cortex, and insensitivity to these regularities in adulthood—demonstrating the importance of perceptual learning in establishing even the core structures on which perception is built. The importance of a mechanism for integrating past visual experience with ‘hard-wired’ circuitry is underscored by what we might consider to be the functional utility of perception, ‘to produce the best current interpretation of the visual scene in light of past experience either of ourselves or of our ancestors’ (Crick and Koch, 1995).

In what follows, we will consider changes to the organization of a sensory object into elements or relatively long-lasting changes to one’s sensitivity to elements to be perceptual changes. While some researchers argue that changes of attention to stimulus elements should be considered pre- or post-perceptual (Pylyshyn, 1999), habitual attention to task relevant features leads to their perceptual sensitization, and affects how the objects are subjectively perceived (see MacPherson, 2012 for a theoretical analysis of some of the evidence for this) as well as perceptual discriminations that one can make (Goldstone, 1994).

1 Individual Differences, Individual Similarities, and the (Im)penetrability of Perception

The role of past experience in shaping the perceptual system suggests the possibility of substantial individual differences in perception. The tailoring of perceptual systems to individuals’ situations raises the concern of unconstrained cognitive penetration. Cognitive penetration of central system beliefs and goals on perception occurs if perceptions can differ by virtue of people having different beliefs or goals.

Some researchers argue that cognitive penetrability does not in fact occur (Pylyshyn, 1999, 2003). Such cognitive penetrability of perception might entail that there is no neutral perceptual ground upon which scientific theory can be verified (Fodor, 1984; see also Chapter VII.2 by Susanna Siegel and Nico Silins), raising the spectre of relativistic theories of meaning. If what we see depends on what we *believe*, then how can scientists from different paradigms use perceptual evidence to adjudicate in favour of one theory or another?

One answer to these concerns has been to claim that perceptual learning is ‘data-driven or task-driven’, and not ‘theory-driven’, meaning that it is the stimulus properties that determine perceptual learning not antecedent beliefs regarding the ‘intentional stance’ of these stimulus properties. This therefore constitutes only an *indirect* cognitive penetrability of perception (via allocation of focal attention to spatial locations in the stimulus). Because in principle, scientists with unequal perceptual learning (but similar perceptual circuitry and scientific training) can receive instructions in how to allocate their attention and therefore obtain an equal perceptual basis, this indirect cognitive penetrability of perception is not problematic for the existence of a theory-neutral perceptual ground, and this dispels relativistic concerns (Raftopoulos, 2001).

However, there is reason to think that perceptual learning is sometimes theory-driven. There *is* an important sense in which certain early visual algorithms (e.g. edge detection) are not directly trainable, and it is certainly not the case that all information influences all processing at all levels. Nevertheless, there are so many ways of guiding perception according to one’s *goals*, that to say that top-down information cannot influence a small stage of early vision seems overly scholastic—because the *functional* consequence is of cognitive, theory-driven, influence on perception. For instance, one might deform the stimulus input by pushing on one’s eye, which changes the character of the stimulus data—a theory-laden influence on the stimulus array, which then produces changes in ‘data-driven’ selective attention. In subsequent sections of this chapter, we consider other evidence on how tasks and goals influence how perceptual systems adapt.

Furthermore, this argument for an across-the-board denial of theory-laden influences on perception seems to presuppose that perceptual systems are essentially identical apart from perceptual learning. However, habitually executed tasks cause not only particular sensory features to be selectively attended to, but also feature creation (Schyns et al., 1998). Feature creation involves creating new perceptual organizations of sensations. For example, brightness and saturation, two components of colour, are psychologically fused for most people in that it is difficult

to attend to one of these components without attending to the other. However, sufficient experience with isolating out colour components can lead colour experts to create psychological dimensions that can be selectively attended to (Burns and Shepp, 1988). Perceptual learning that requires allocating attention to previously created features can count as theory-driven perceptual learning given the importance of goals and tasks on the original construction of perceptual features. It is important to distinguish between the role of goals in shaping perception over a long time course and the role of goals in shaping one's moment-to-moment perception without permanently changing a perceptual system. Here we are focusing on the former kind of goal-dependence. Given the developmental evidence, it seems likely that much of perceptual learning builds upon previously created perceptual building blocks. Pylyshyn (1999, 2003) grants that cognition can shape the operation of vision into certain 'compiled transducers', computational routines that become automatized through practice and eventually become encapsulated as part of the early vision system itself. Although this diachronic change does not count as direct cognitive penetrability of perception, it entails that people with different experiences could have fundamentally different perceptual systems. This, in turn, is problematic for claims that equivalent training can equalize perceptual differences among scientists, and consequently problematic for the existence of a theory-neutral perceptual ground.

Given the ubiquity of top-down influences on perception (even on ostensibly 'early' visual processes, e.g. Chen and Zhou, 2011), some have questioned the value of drawing a firm boundary between perception and cognition at all (Goldstone and Barsalou, 1998; Barsalou, 1999). While we will not defend the radical claim of continuity between perception and cognition here, gross anatomical considerations do indicate that when one brain region X sends feed-forward information to region Y, there are typically dense recurrent connections from Y to X.

A theory-neutral foundation for perceptual processes is unlikely to exist in practice, either at a low level or a high one because, for example, novices and bird experts perceive the same birds differently, even given the same-retinal input (Gauthier et al., 2010). However, this does not preclude the lack of such a shared basis *in principle*, given common training experiences. While there are striking differences in individuals' perceptual processes due to their experiences, the process by which perceptual systems change with experience is largely shared across individuals. In other words, providing *equivalent* perceptual learning opportunities generally produces equivalent perceptual changes. (See Olshausen and Field, 1996, for one model, though on an evolutionary time scale.) This shift from a first-order to second-order invariant across people grants that equivalent perceptual learning opportunities involve more than just data-driven perceptual learning. By questioning a strict boundary between perception and conception, we can suppose that with shared communication, a shared familiar world of physical objects, shared scientific and perceptual training, and shared developmental trajectories, the shared perceptual-cognitive basis is probably 'neutral enough' for successful scientific communication (Swayer, 2003).

2 Evidence and Rationale for Early Loci of Perceptual Change

During learning, changes to the human perceptual system occur at multiple stages in the information processing stream. There is general consensus that changes to earlier, more peripheral stages of processing are more unambiguously identifiable as 'genuine' cases of perceptual learning compared to later, more central stages. This belief hinges on the assumption that later stages of information processing in the brain may not reflect perceptual processes that are uncontaminated by context and experience. There is evidence, however, that learning and context influence even relatively early stages of perceptual processing. This evidence is based on both neuroscience and functional behaviour.

Neurological evidence for early changes to perception

Neurologically speaking, changes to early-to-middle stages of visual processing have been implicated in the development of expertise. Electrophysiological recordings of dog and bird experts show enhanced electrical activity at 164 milliseconds after the presentation of dog or bird pictures (Gauthier et al., 2010). Practice in discriminating small motions in different directions significantly alters electrical brain potentials that occur within 100 milliseconds of the stimulus onset (Fahle, 1994). These electrical changes are centred over the part of the visual cortex primarily responsible for motion perception (the medial temporal visual area, MT), suggesting plasticity in early visual processing. Furmanski et al. (2004) used functional magnetic resonance imaging (fMRI) to measure brain activity before and after one month of practice detecting hard-to-see oriented line gratings. Training

increased V1 response for the practised orientation relative to the other orientations, and the magnitude of V1 changes were correlated with detection performance. Bao et al. (2010) trained human subjects for one month to detect a diagonal grating, and found EEG differences in V1 for trained versus untrained orientations within 50–70 milliseconds after the onset of the stimulus. The rapidity of the EEG difference, combined with the demanding nature of the primary behavioural task during testing make it unlikely that the earliest EEG differences were mediated by top-down feedback from higher cortical levels. In the somewhat later visual area V4, single-cell recording studies in monkeys have shown activity changes of cells in early visual cortex (Yang and Maunsell, 2004). Individual neurons with receptive fields overlapping the trained location of a line orientation discrimination developed stronger responses, and more narrow tuning, to the particular trained orientation, compared with neurons with receptive fields that did not fall on the trained location.

In the auditory modality, Weinberger (1993) describes evidence that cells in the primary auditory cortex become tuned to the frequency of often-repeated tones. Training in a selective attention task produces differential responses as early as the cochlea (Puel et al., 1988). This amazing degree of top-down modulation of a peripheral neural system is mediated by descending pathways of neurons that project from the auditory cortex all the way back to olivocochlear neurons, which directly project to outer hair cells within the cochlea—an impressively peripheral locus of modulation.

More generally, we do not find it particularly productive to ask the question of ‘How early does perception change due to learning?’ because it is clear that prior learning affects sensory processing even before sensory processing begins. In particular, learning influences how objects will impinge upon our sensory organs. In many cases, perceptual learning involves acquiring new procedures for actively probing one’s environment (Gibson, 1969), such as learning procedures for efficiently scanning the edges of an object (Salapatek and Kessen, 1974). The result is that adults look at objects differently than children, and experts look at objects differently than novices; and since each fixates objects differently, the visual patterns that fall on an observer’s retina vary with experience. Perceptual changes are found at many different neural loci and a general rule seems to be that earlier brain regions are implicated in finer, more detailed perceptual training tasks (Ahissar and Hochstein, 1997).

The claim for widespread neural plasticity in brain regions related to perception should not be interpreted as an argument for the equipotentiality of brain regions for implementing modifications to perception. Evidence for plasticity at the earliest visual processing area of the cortex, V1, remains controversial (Crist et al., 2001; Kourtzi and DiCarlo, 2006). Some of the observed activity pattern differences in V1 may be attributable to top-down influences after a first forward sweep of activity has passed. However, the very presence of large recurrent connections from more central to more peripheral brain regions attests to the evolutionary importance of tailoring input representations to one’s tasks. Properties of V1 cells depend on the perceptual task being performed and experience, in the sense that neurons respond differently to identical visual patterns under different discrimination tasks and with different experiences. Moreover, these top-down influences are seen from the onset of neural response to a stimulus (Li et al., 2004). The perceptual change, thus, is early both in the information processing stream of the brain and chronometrically.

One common source of evidence for an early neural locus for perceptual learning has been observations of surprisingly limited transfer of learning. Training on simple visual discriminations often does not transfer to different eyes, to different tasks, to different spatial locations (Shiu and Pashler, 1992), or to different viewing distances (Huang et al., 2011). The customary interpretation of these narrow degrees of generalization is that early perceptual detectors tend to have narrow and small receptive fields, and that downstream detectors have larger receptive fields. So, if narrow generalization is found, it is likely to be driven by the earlier detectors. In truth, changes to perceptual systems are found at multiple stages during perceptual processing (Ahissar and Hochstein, 1997), with some kinds of perceptual learning being mediated by more general and strategic changes. Accordingly, perceptual learning is not as restricted as one might suspect according to a ‘data-driven’ account in which perceptual change is accomplished by a perceptual process being passively imprinted upon by environmental stimuli. A role for strategic adaptation is also suggested by the finding that limited generalization of perceptual learning can be greatly reduced by giving observers a small amount of training at the transferred direction or location (Xiao et al., 2008).

Functional evidence for early changes to perception

Parallel to neurological evidence for loci of perceptual changes that accompany experience, there are functional, behavioural sources of evidence indicating changes to early stages of perceptual processing. Experience often exerts an influence before other putatively early perceptual processes have been completed. For example, subjective experience exerts an influence on colour perception before the perceptual stage that creates colour after-images has completed its processing (Moscovici and Personnaz, 1991). As a second example, Peterson and Gibson (1994) found that the organization of a scene into figure and ground is influenced by the visual familiarity of the contours. Their participants made judgements about whether a visual form was a figure or ground. Familiar forms were more likely to be judged to be figures than unfamiliar forms. This effect was not found when the familiarity of the objects was eliminated by flipping the scenes upside down. Interpreting the familiar region as a figure was found even when the unfamiliar regions had the strong Gestalt organization cue of symmetry. Peterson and Lampignano (2003) found direct evidence that the acquired familiarity of a shape successfully competes against Gestalt cues such as partial closure to determine the organization of a scene into figure and ground.

Perceptual organizations that are natural according to Gestalt laws of perception can be overcome in favour of perceptual organizations that involve familiarized materials. This indicates that training can influence relatively early stages of the information processing stream. In work on object-based attention, Behrmann, Zemel, and Mozer (1998) found that judgements about whether two parts had the same number of humps were faster when the two parts belonged to the same object rather than different objects, with objecthood being based on standard Gestalt laws of organization. However, follow-up work found an influence of experience on subsequent part comparisons. Two fragments were interpreted as belonging to the same object if they had co-occurred many times in a single shape (Zemel et al., 2002). Object fragments that are not naturally grouped together because they do not follow the Gestalt law of good continuation can nonetheless be perceptually joined if participants are familiarized with an object that unifies the fragments.

Another approach to identifying the functional locus of perceptual changes is to observe the time course of the use of particular types of information. For example, on the basis of priming evidence, Sekuler, Palmer, and Flynn (1994) argue that knowledge about what an occluded object would look like if it were completed influences processing after as little as 150 milliseconds. Dog and bird experts reveal significantly enhanced N170¹ electrophysiological responses when categorizing objects within their domain of expertise relative to objects outside of this domain (Tanaka and Curran, 2001). These influences are sufficiently early to typically be counted as perceptual processing.

Reasons why early perceptual learning occurs

The plasticity of early perceptual processes may seem counterproductive. Don't we want perception to act as a source of information that is uncontaminated by a perceiver's beliefs and history? To the extent that we all live in the same physical world, shouldn't we all be equipped with the same perceptual apparatus? There is something right about this intuition. Our early perceptual processes should change slowly and conservatively, because they are the bedrock for all subsequent processes. A change early in the brain will have ripples of influence downstream.

The answer, though, to why it may still be a good idea to adapt early perceptual processes to experiences is that flexibility is beneficial when the world is variable. If everyone were confronted with the same environment, and this environment remained unchanged millennium after millennium, then perceptual systems could become hard-wired for this particular environment. These perceptual systems would be efficient because they are specifically tuned to the unchanging environment. Some environmental factors, such as colour characteristics of sunlight, the position of the horizon, and the change in appearance that an approaching object makes, have all been mostly stable over the time that the human visual system has developed.

However, if we look more closely, there is an important sense in which different people face different environments. To a large extent, a person's environment consists of animals, people, and things made by people. Animals and people have been designed by evolution to show variability, and artefacts vary widely across cultures. Evolutionary pressures may have been able to build a perceptual system that is generally adept at processing faces (Gauthier et al., 2010), but they could not have hardwired a neural system that was adept at processing a particular face, such as Barack Obama's, for the simple reason that there is too much generational variability among faces. Individual faces do not last from generation to generation, and so people's ability to recognize

specific, highly familiar faces cannot be hardwired. Rather, what is hardwired is the ability to develop perceptual systems tuned to particular faces. Variability is apparent over only slightly longer periods for artefacts, words, ecological environments, and animal appearances. Thus, we can be virtually positive that tools show too much variability over time for there to be a hardwired detector for hammers. Words and languages vary too much for there to be a hardwired detector for the written letter 'A'. Biological organisms are too geographically diverse for people to have formed a hardwired cow-detector. When environmental variability is high, the best strategy for an organism is to develop a general perceptual system that can adapt to its local conditions.

There is an even deeper sense in which people face different environments. People find themselves in different worlds because they choose to specialize. English-speaking people become specialized at hearing and seeing English words. People familiar with a particular race become specialized at recognizing faces from that race (Levin, 2000). Experts at wine tasting, chick sexing, X-ray diagnosing, identical twin identifying, and baseball pitch judging all have unique perceptual capabilities because of the tasks they perform. Experts typically have highly specialized skills, many of which are perceptual (Sowden et al., 2000). Moreover, the above examples of word and face recognition suggest that every person has domains in which they show expertise. Even if all people confronted the same world initially, they would create distinctive communities with unique languages, artefacts, and objects of importance. One's social niche will depend on many factors including proclivity, community needs, local support, random accidents, and positive feedback loops. Thus, it is again advisable to build a perceptual system with the flexibility needed to support any one of a large number of niches.

It is worth emphasizing that people do not have to strategically, consciously tune their perceptual systems to support these tasks that are specific to an individual's needs and niches. Instead, the perceptual machinery itself, and the brain areas it projects to, can often automatically accomplish the necessary tuning. Likewise, perceptual tuning can occur independently of a perceiver's antecedent beliefs. In fact, once a tailored perceptual system is firmly in place, it is impossible to revert it to an untailored state even if one is motivated to do so by momentary beliefs or desires.

3 Mechanisms of Perceptual Change

The above neurophysiological and functional evidence makes a good case for relatively early changes to perception due to experience, particularly if one includes changes to eye fixations that change perceptual differences even before the retina. This case having been briefly made, the discussion that follows focuses not on the loci of changes, but rather on the mechanisms that underlie some of these changes. These mechanisms will be described at a functional rather than physiological level. Even though neurological details are known in some cases, a functional level of description is appropriate for unifying accounts of human and computational learning (Goldstone, 2003), and for understanding the kinds of changes that may benefit a perceiver.

Tuned attention

A person can dynamically shift their attention to different stimulus features depending on their perceived importance. One way that perception becomes adapted to tasks and environments is by increasing the attention paid to perceptual features that are important, and/or by decreasing attention to irrelevant dimensions and features. This mechanism can be distinguished from a more passive 'imprinting' process in which functional perceptual detectors are developed that are specialized for stimuli or parts of a stimuli. In computational learning terms, imprinting is an unsupervised learning mechanism in that there is no need for a parent, teacher, or programmer to tell the learner what a stimulus is called or what parts are important. However, in attentional tuning, supervision, including self-supervision, is key to learning. In distinguishing ripe from unripe mangoes, colour must be attended to, but to distinguish books from magazines, colour is not useful. An observer equipped with the ability to tune their attention to different object dimensions would be able to learn both of these categorizations.

The world is structured such that different dimensions are important for different life-relevant categorizations. As such, it comes as little surprise that most successful theories of categorization and learning incorporate selective attention. To take Nosofsky's (1986) Generalized Context Model of categorization as an example, the categorization of an object depends on its similarity to previously stored category members in a multidimensional space. Critically, psychological distances between objects are compressed and expanded along dimensions in this

space depending on the categorization required. Distances between objects on relevant dimensions are expanded while irrelevant dimensions are compressed. For example, Nosofsky finds that if participants are given a categorization where the angle of a line embedded in a circular form is relevant while the size of the circular form is irrelevant, then distances between objects on the angle dimension are increased and distances on the size dimension are decreased.

In the Generalized Context Model, entire perceptual dimensions like length, brightness, and orientation are psychologically stretched or shrunk. However, specific regions within a perceptual dimension can also be selectively attended. This capacity is important for driving the phenomenon of Categorical Perception (CP) (Goldstone and Hendrickson, 2010; also see Raffman). In CP, our perceptions are adapted such that differences between objects that belong to different categories are accentuated, and differences between objects that fall into the same category are deemphasized. That is, our perceptual systems transform relatively linear sensory signals into relatively non-linear internal representations. The extreme case of this kind of non-linear transformation is a step function by which increases to a sensory signal have no effect on perception until the signal reaches a certain threshold. At that threshold, perception qualitatively and suddenly changes. During the flat portion of the staircase function, different input signals have equivalent effects. This flat response profile for a range of stimuli provides a mechanism that grounds equivalence classes—for treating distinguishable stimuli as equivalent. Equivalence classes, in turn, provide us with the beginnings of symbolic thought—quasi-discrete responses that are reliably generated when stimuli within a range are presented or contemplated.

The underpinning for the non-linear perceptions of CP is region-specific attention tuning. For example, in Liberman, Harris, Hoffman, and Griffith's (1957) seminal research on speech perception, a continuum of equally spaced consonant-vowel syllables with endpoints reliably identified as /be/ and /ge/ was generated. At a specific point along this continuum observers rapidly shift from hearing the sound as a /be/ to hearing it as /de/. In addition to giving participants an identification task, participants were also given an ABX discrimination task. In this task, observers listened to three sounds—A followed by B followed by X—and indicated whether X was identical to A or B. Observers performed the task more accurately when syllables A and B belonged to different phonemic categories, as indicated by their identification probabilities, than when they were variants of the same category, even when physical differences were equated.

Although some cases of CP may be innate, there is also strong evidence that at least some cases involve learning. Using laboratory-created, speech-like stimuli that were assigned to different categories based on their labels, Lane (1964) found CP effects despite a lack of correspondence between the trained categories and naturally occurring language categories. In general, a sound difference that crosses the boundary between phonemes in a language will be more discriminable to speakers of that language than to speakers of a language in which the sound difference does not cross phonemic boundaries (Repp, 1984). Goldstone (1994) showed that when arbitrary new visual categorizations are made, they alter participants' same/different judgements. Discriminations along categorization-relevant dimensions were improved, and this improvement was greatest at the boundary between the categories.

The attention tuning phenomena reviewed above indicate an important interplay between humans' higher-level conceptual systems and their lower-level perceptual systems. Traditional information flow diagrams in cognitive science typically draw a clean division between perceptual and conceptual systems, with information moving only from perception to the conceptual system. The common occurrence of attentional tuning indicates permeability and bidirectional influence between these systems. We humans do not simply base our categories on the outputs of perceptual systems independent of feedback. Instead, our perceptual systems become customized to the task-useful categories that we acquire. We are not optimistic that a clean dividing line between perception and attention can be drawn because (1) with training, the sensory encoding for attended objects becomes richer and more differentiated than for unattended objects; (2) with training, attention to objects becomes automatically deployed and difficult to strategically control; and (3) fast and widely prevalent recurrent connections from higher to lower cortical regions makes it difficult, sometimes impossible, to identify a 'forward-volley' stage of sensory processing that is uninfluenced by attention.

Unitization

Beyond simply tuning attention to existing perceptual dimensions or regions of a dimension, there are two

perceptual learning mechanisms that create new 'building blocks', new perceptual units that can be attended to, searched for, and combined together to form new percepts. One of these mechanisms is unitization, according to which single functional units are constructed that are triggered when a familiar complex configuration arises. Unitization is the perceptual equivalent of the memory-based phenomenon of chunking. In memory, even if we can only store 7 +/- 2 items in our short-term memory, we can learn to store increasing amounts of information by increasing the size of each of the items (chunks) in memory. We can easily remember the 27 letters 'M O N T U E W E D F B I C I A K G B C B S N B C A B C' even though there are far more than 7 letters by combining the letters to create acronyms such as MON, FBI, and CBS. Chunking at a still higher level, we can remember that the acronyms form three categories: days of the week, secretive government organizations, and television broadcast channels. Similarly, our perceptual systems can build new perceptual chunks to encode as single units what would otherwise be complex visual patterns. Whereas for chunks in memory elements are combined together because they make semantic sense, perceptual units are formed because they can be seen as coherent perceptual objects, obeying, for example, the Gestalt laws of proximity, continuity, and closure.

Cattell (1886) invoked the notion of perceptual unitization to account for the advantage that he found for tachistoscopically presented words relative to non-words. Unitization has also been posited in the field of attention, where researchers have claimed that shape components of often-presented stimuli become processed as a single functional unit with practice. Shiffrin and Lightfoot (1997) report evidence from the slopes relating the number of distracter elements to response time in a feature search task. When participants learned a conjunctive search task in which three line segments were needed to distinguish the target from distracters, impressive and prolonged decreases in search slopes were observed over 20 hour-long sessions. These prolonged decreases were not observed for a simple search task requiring attention to only one component. The authors concluded that conjunctive training leads to the unitization of the set of diagnostic line segments, resulting in fewer required comparisons.

Gauthier and Tarr (1997; see also Gauthier et al., 1998) found that prolonged experience with a novel object leads to a configural representation of it that combines all of its parts into a single, viewpoint specific, functional unit. Their evidence for such a representation is that recognition of these familiarized objects improved considerably with practice, and was much more efficient when the object was in its customary upright form rather than inverted. Two more recently developed methods for detecting processing based upon single functional units for an entire complex object, also called 'holistic' processing, involve 'whole-part' and 'composite' tasks.

The logic of the whole-part paradigm is that if a stimulus element is being processed as part of a larger perceptual unit, then recognition of the part should be less efficient than recognition of the entire unit. For example, Tanaka and Farah (1993) show subjects a whole face named 'John'. After a brief delay, they are either asked to identify which of two noses in isolation, say John's or Kevin's, was the previously displayed nose. On other trials, they are asked to say which of two faces was John's face; one of the faces is indeed John's face, and the other face is exactly the same as John's face except that John's nose has been replaced by Kevin's nose. Even though both tasks ostensibly require selecting the face with John's nose rather than Kevin's nose, subjects are better in the whole face than the isolated part condition. This result is reminiscent of earlier demonstrated word-superiority and object-superiority effects showing that perception of parts in the context of larger familiar objects is better than perception of isolated parts (Wheeler, 1970), and has been explained in terms of the whole face being a functional unit of perception.


In the composite task, a composite face is formed by combining the top and bottom parts of possibly different individual faces. Subjects are tasked with responding to only the top or bottom part of the face while trying to ignore the other half (Carey and Diamond, 1977). In fact, subjects find it difficult to ignore the irrelevant half. For example, if the top part of a composite belongs to John but the bottom belongs to Kevin, then it is difficult for subjects to respond 'John' to the top part, compared to a situation in which both halves belong to John. The distracting influence of the irrelevant part is reduced or eliminated if the halves are misaligned—if they no longer meet to create to form an apparently coherent face (Gauthier and Bukach, 2007).

From the evidence considered thus far from both the whole-part and composite paradigms, it might simply be concluded that humans have been evolutionarily wired such that whole faces are the unit of perception. It is, indeed, rare to stumble across isolated noses, and when one does, they are not usually social objects of importance. However, the units implicated by these tasks are not only faces, but are more generally objects with

which an observer has had prolonged experience (Gauthier et al., 2010). Empirical evidence for perceptual units has been found for objects as diverse as birds, words, grids of lines, random wire structures, fingerprints, artificial blobs, and three-dimensional creatures made from simple geometric components, as long as these objects have been familiarized over at least somewhat protracted time courses. Units like these may have been created from more elemental parts, but once they have been formed, they renounce their origins. Once created, the units operate such that perceptions are no longer experienced in terms of the parts, but rather in terms of the whole.

Attribute differentiation

If unitization creates large, complex units out of elemental parts, then attribute differentiation begins with a complex perception that fuses together multiple components, and then develops separate percepts for the individual components. In this sense, unitization and attribute differentiation are complementary mechanisms of perceptual learning. However, in another sense they are flip sides of the same coin—they both involve creating perceptual units that are tailored to one's tasks and experiences.



Click to view larger

Figure 01: Stimuli from Goldstone and Steyvers (2001). Every face in the 4×4 grid is formed by combining in equal parts its value on Dimensions A and B. These dimensions, in turn, are formed by varying the proportion of two randomly chosen faces. Increasing values of Dimension A correspond to simultaneously decreasing the amount of Face 1 and increasing the amount of Face 2. Similarly, increasing values of Dimension B correspond to decreasing the amount of Face 3 as the amount of Face 4 increases.

To understand the role of attribute differentiation, it is useful to return to the first described mechanism—tuned attention. Attention tuning requires the ability to selectively attend to perceptual attributes that have already been psychologically isolated. That is, it is only possible for an observer to attend to the brightness of a shape and ignore its size if the attribute of brightness has been isolated and separated from size. What happens if two attributes have *not* yet been isolated? Saturation (a psychological dimension related to the amount of white/black added to a colour) and brightness (a psychological dimension related to the amount of luminance energy emitted by a colour) are two such attributes. They are fused together to form an overall impression of colour. Ordinarily, people cannot selectively attend to just the saturation of a colour, ignoring its brightness, or vice versa (Garner, 1976). However, there is evidence that people can learn to selectively attend to attributes. If subjects are given training in which saturation is relevant for a categorization and brightness is irrelevant, they can learn to form this categorization, and when they do, subjects subsequently find it easier to distinguish between objects on the basis of saturation rather than brightness differences (Goldstone, 1994). A particularly efficient way to separate the attributes of saturation from brightness is to repeatedly alternate training where saturation is task-relevant with training where brightness is task-relevant, with the end result that either of the attributes can be selectively attended while ignoring the other (Goldstone and Steyvers, 2001). Colour experts such as artists or vision scientists are better able to selectively attend to the component dimensions of colour than are novices (Burns and

Shepp, 1988). Even arbitrary dimensions such as those created by morphing between randomly selected faces can be isolated with training that repeatedly makes one, then the other, relevant for a categorization (Goldstone and Steyvers, 2001). The arbitrary dimensions shown in Figure 01 do not start off being perceptually isolated for an observer, but can come to be selectively attended with practice that requires their isolation.

There is developmental evidence that attributes that are easily isolated by adults, such as the brightness and size of a square, are treated as fused together for 4-year old children (Smith and Kemler, 1978). It is relatively difficult for children to decide whether two objects are identical on a particular attribute, but relatively easy for them to decide whether they are similar across many attributes. For example, children seem to be distracted by shape differences when they are instructed to make comparisons based on colour. From a functional perspective, we interpret this intrusion of one dimension when trying to respond to another dimension as indicating relatively fused perceptual encodings in children. This interpretation is supported by neurophysiological evidence that children's cortical sensory areas are not as specialized for specific modalities as are adults' (Maurer and Mondloch, 2004). Children's brains have cross-modal connections between auditory, visual, and tactile sensory areas that are subsequently pruned with experience (Spector and Maurer, 2009).

Attribute differentiation is not equivalent to tuned attention because prior to differentiation training that separates two attributes, it might have been impossible for an observer to selectively attend to either of these attributes while ignoring the other. After differentiation training, there is a longer-term ability of the observer to switch their attention to *either* of the two attributes. If selective attention to an attribute can be understood as learning to weight that attribute heavily for a judgement, then attribute differentiation can be understood as *learning to learn* how to weight an attribute. For example, once an individual has perceptually differentiated the saturation of a colour from its brightness, then s/he can quickly learn a discrimination based on either attending saturation or brightness, quickly learning to switch attention to whichever dimension is relevant at a given time. The extent of dimension differentiation determines the efficiency with which newly relevant dimensions can be selectively attended.

Attribute differentiation is of theoretical interest because even the possibility of it may be denied. One might suppose that if two attributes are fused together at some point in perceptual processing, then they can never be split apart later. Once yellow and blue watercolour paint have been mixed, they cannot be unmixed. Fortunately, there are computational models that explain how attribute differentiation mechanisms might operate. Competitive learning neural networks differentiate inputs into categories by specializing detectors to respond to classes of inputs. Random detectors that are slightly more similar to an input than other detectors will learn to adapt themselves toward the input and will inhibit other detectors from doing so (Rumelhart and Zipser, 1985). A model that extended this mechanism to sorting object parts into detectors, when presented with an original set of training objects, was able to discover part-based building blocks that could be recombined to recreate the original training objects (Goldstone, 2003). Another learning system shows similar functional behaviour by using Bayesian methods (Austerweil and Griffiths, 2011). In short, advances in machine learning provide existence proofs of mechanisms for dimension differentiation. Computationally speaking, green paint can be separated into its yellow and blue components if one has not only a single sample of green, but several colour samples with different proportions of yellow and blue.

Summary of mechanisms of perceptual learning

The perceptual representations that result from distal objects are influenced by experiences with the objects. By attentional tuning, pre-existing perceptual attributes are sensitized or desensitized. By unitization, originally separated parts of an object are combined into unified and coherent perceptual wholes. Once constructed, the unit can be efficiently recognized and has properties similar to an image-like template. The opposite process, attribute differentiation, can also occur, separating originally integrated percepts into psychologically differentiated dimensions or parts. Rather than viewing unitization and differentiation as contradictory, they are best viewed as aspects of the same process that bundles stimulus components together if they diagnostically co-occur, and separates these bundles from other statistically independent bundles. Under this conception, learning a perceptual organization consists in learning how to carve a stimulus into useful components. These empirical phenomena, and their associated computational models, strongly suggest that perceptual learning is affected by our concepts. To be sure, our perceptions also ground our concepts, but interestingly, they provide a *better* grounding for our concepts because they are flexibly altered by these concepts. Like a mattress that provides support by

conforming to the sleeping body that lies on it, our perceptions support our concepts by conforming to them.

4 Modifying Perception to Get What One Wants

There is some evidence that we change our perceptual systems even when this is not what we intend to do. For example, mere exposure to perceptual information that is irrelevant to a task can suffice for improving an observer's perceptual sensitivity to this information (Watanabe et al., 2001). Moreover, this kind of perceptual learning can occur even when the information is in the visual periphery and subthreshold, that is, below the threshold for conscious detection.

However, complimenting these results are other results showing that what is learned and *how* efficiently it is learned depends on the observer's task and goal. Even when sensitivity to a line orientation appears to have a relatively early locus of change, in that it does not transfer strongly across eyes or visual regions, it nonetheless depends on the observer's goals (Shiu and Pashler, 1992). Perceptual sensitization to a line orientation is much more robust when it is relevant for the task than when it is irrelevant. When observers are given the same stimuli in two conditions, but are required to make fine, subordinate-level categorizations in one condition but coarser, basic-level categorizations in the other, then greater selectivity of cortical regions implicated in object processing were found in the former condition (Gillebert et al., 2008). Finally, much of the evidence for categorical perception (see Chapter VI.1 by Diana Raffman) indicates that perceptual discriminations are easier to make at boundaries between important categories for an observer, such as between a /p/ and /b/ phoneme that would be important for distinguishing 'pats' from 'bats'. Evidence from training studies and cross-linguistic comparisons indicates that it is not just perceptual sensitivities that are driving the categories, but rather the important categories are also driving perceptual sensitivities (Goldstone and Hendrickson, 2010). All of these studies show that we get selectively better at making perceptual discriminations just where we *should* get better in order to help us do what we want and need to do.

This observation is particularly relevant to the question of cognitive penetrability (see Chapter VII.2 by Susanna Siegel & Nico Silins). What we see is influenced by repeatedly possessed goals that are close to our biologically determined needs and drives. This is different from the claim that our in-the-moment goals influence our perceptions, which we also believe happens but requires a different assembly of empirical justifications. Our current claim is that the habitual discriminations and categorizations that we make influence our perceptual abilities in a particular functional direction—we selectively improve our perceptual abilities so that the tasks that we need to perform are performed better. In the previous sentence, there are two interpretations of 'so that'. By one interpretation, 'so that' means 'with the intention that', implying that we strategically and explicitly alter our perceptual abilities. By the second interpretation, 'so that' means 'with the end result that', implying that our perceptual abilities are altered naturally through an automatic, non-conscious process. We will return later to the former interpretation, but our main interest lies with the latter interpretation, which has a strong evidentiary basis and is theoretically important as well.

The alterability of our perceptual systems in a personally useful direction without our explicit intention to do so is simple but powerful. It is simple in that it requires no more sophisticated a mechanism than random variation plus selection. The effective strengths of neuronal connections are constantly varying. If a change causes important discriminations to be made with increasing efficiency, then the change tends to be preserved and extended. If not, the change will not be made permanent. Dopamine plays a key role during reinforcement learning in which external or internal rewards drive a learning consolidation process. For example, in a motion discrimination task, incorrectly predicting a reward guides changes in connections involving the primary brain area dedicated for motion processing (MT), by selectively strengthening the connections from the most sensitive neurons in the sensory population (Law and Gold, 2009). There may be other more goal-directed processes of neuronal change, but simple random variation with reinforcement that is potentially internally generated suffices to systematically improve perceptual systems at a longer time scale. This systematicity is the key to the power of these changes. Even without opening up the black box of a perceptual module (assuming such a beast to exist), it is possible to make these modules reliably improve their performance much more often than they degenerate, as long as the observer, the environment, or a teacher of some sort can provide feedback on whether the observer is doing better or worse after a change than before.

Returning to the question of strategic alterations to perceptual systems, there do appear to be cases in which people purposefully 'hack' their perceptual systems in order to facilitate performance. Through trial and error we learn that we can create a sharper image of an object by arranging our fingers so as to create a small aperture near our eye. We learn that cupping our hands behind our ears allows us to hear better, whereas clamping our jaw tight makes our ears less sensitive to noise. Altering one's physical interactions with the world is a major part of Gibson's (1979) theory of 'active perception'. At a second-order level, we learn over time how to improve our own learning. If we are trying to become an expert at identifying birds, we have learned to facilitate this by studying photos of birds repeatedly. Depending on our sophistication and knowledge of the learning process, we may strategically present birds in more caricatured forms such as idealized drawings rather than photographs, in different viewpoints, with time in between presentations, and with different species intermixed. These are all strategic actions that improve perceptual learning, and they can all be accomplished without requiring cognitive penetration that involves our direct alteration of neuronal processes underlying perception.

There are many intermediate cases in which it is difficult to tell whether we intend to change our perceptual systems or they just naturally change. For example, work in our laboratory has shown that we change our visual processing system so that it allows us to reason in a formally sanctioned way more efficiently than we would have been able to without the change. In particular, we have studied the ways in which we 'rig up' our visual processing to facilitate mathematical reasoning (Landy and Goldstone, 2007; Goldstone et al., 2010). In algebra, multiplication has a higher order of precedence, than addition with '2 + 4 × 5' equalling 22, not 40. Experiments show that one way that we come to produce the correct solution efficiently is by automatically directing our own attention to the '×' operator instead of the '+' operator. Our first eye movements are toward the '×' operator, and the '×' operator distracts our attention from a task requiring attention to the '+' operator more so than vice versa (Goldstone et al., 2010). We learn to solve algebra problems by creating visual operations like moving a notation element from the left to the right side, marking notational elements as processed, cancelling matching notational elements, and directing attention to mathematical groups. Although these operations are 'merely' perceptions and actions, these physical operations can be adjusted and tailored so that they conform to formally sanctionable operations.

If evolution has done a good job of designing perceptual learning processes, then it should be the case that naturally occurring changes to perception will often mimic the changes that would have been produced had we strategically orchestrated the changes. Thus, for the case of developing perceptual routines to help us solve formal math problems, we may consciously be training our perceptual systems to 'do the right thing', or we may just be relying on automatic processes of fluency, habit, and trainable perceptual grouping to help us. General-purpose neural processes of reinforcement learning, associative learning, and hierarchical action chunking help us to perform increasingly sophisticated activities by shunting activities that originally required strategic planning and central executive processes over to fast perception-action processes. Whereas some philosophers have argued that a hallmark of an advanced science is that it no longer requires notions of perceptual resemblance as the basis for its categories (Quine, 1977), the argument from perceptual learning is that our perceptual systems are not doomed to use untutored and potential misleading perceptual systems. At first sight, $\frac{X+4}{3+4}$ may resemble $\frac{X*4}{3*4}$ but to the perceptually trained student of math, only the 4s in the second case *look like* they can be cancelled. At first sight, the marsupial Tasmanian wolf may resemble its quite distant evolutionary cousin, the placental grey wolf, but to a trained biologist, their teeth, jaws, ears, and nasal cavities are quite obviously distinct. If one wishes to become cognitively more sophisticated, an alternative to trumping perception is to train one's perceptual systems instead.

Acknowledgements

The authors wish to express thanks to David Landy, Mohan Matthen, Susanna Siegel, and Linda Smith for helpful suggestions on this work. This work was funded by National Science Foundation REESE grant 0910218 and Department of Education IES grant R305A1100060. More information about the laboratory can be found at <<http://cognitn.psych.indiana.edu>>. Correspondence concerning this article should be addressed to Robert Goldstone, Psychological and Brain Sciences, Indiana University, Bloomington, IN, 47405. Email: rgoldsto@indiana.edu.

References

- Ahissar, M. and Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature*, 387, 401–406.
- Austerweil, J. and Griffiths, T. L. (2011). A rational model of the effects of distributional information of feature learning. *Cognitive Psychology*, 63, 173–209.
- Bao, M., Yang, L., Rios, C., and Engel, S. A. (2010). Perceptual learning increases the strength of the earliest signals in visual cortex. *Journal of Neuroscience*, 30, 15080–15084.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Science*, 22, 577–660.
- Behrmann, M., Zemel, R. S., and Mozer, M. C. (1998). Object-based attention and occlusion: Evidence from normal participants and a computational model. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 1011–1036.
- Blakemore, C. and Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature*, 228, 477–478.
- Burns, B. and Shepp, B. E. (1988). Dimensional interactions and the structure of psychological space: The representation of hue, saturation, and brightness. *Perception and Psychophysics*, 43, 494–507.
- Carey, S. and Diamond, R. (1977). From piecemeal to configurational representation of faces. *Science*, 195(4275), 312–314.
- Cattell, J. M. (1886). The time it takes to see and name objects. *Mind*, 11, 63–65.
- Chen, L. and Zhou, X. (2011). Visual apparent motion can be modulated by task-irrelevant lexical information. *Attention, Perception, & Psychophysics*, 73, 1010–1015.
- Churchland, P. (1989). The anti-realist epistemology of Van-Fraassen's The Scientific Image. In B. A. Brody and R. E. Grandy (eds), *Readings in the Philosophy of Science*, Englewood Cliffs, N.J.:Prentice Hall.
- Crick, F. and Koch, C. (1995). Are we aware of neural activity in primary visual cortex? *Nature* 375, 121–123.
- Crist, R. E., Li, W., and Gilbert, C. D. (2001). Learning to see: experience and attention in primary visual cortex. *Nature Neuroscience*, 4, 519–525.
- Fahle, M. (1994). Human pattern recognition: Parallel processing and perceptual learning. *Perception*, 23, 411–427.
- Fodor, J. (1984). Observation reconsidered. *Philosophy of Science*, 51, 23–43.
- Furmanski, C. S., Schluppeck, D., and Engel, S. A. (2004). Learning strengthens the response of primary visual cortex to simple patterns. *Current Biology*, 14, 573–578.
- Garner, W. R. (1976). Interaction of stimulus dimensions in concept and choice processes. *Cognitive Psychology*, 8, 98–123.
- Gauthier, I. and Bukach, C. (2007). Should we reject the expertise hypothesis? *Cognition*, 103, 322–330.
- Gauthier, I. and Tarr, M. J. (1997). "Becoming a" Greeble" expert: exploring mechanisms for face recognition." *Vision research* 37.12,1673–1682.
- Gauthier, I. Tarr, M. J., and Bubbs, D. (eds) (2010). *Perceptual Expertise: Bridging brain and behavior*. Oxford: Oxford University Press.
- Gauthier, I., Williams, P., Tarr, M.J., and Tanaka, J.,(1998). Training 'greeble' experts: a framework for studying expert object recognition processes. *Vision research*, 38(15), 2401–2428.
- Ghahramani, Z. (1995). Factorial learning and the EM algorithm. In G. Tesauro, D. S. Touretzky, and T. K. Leen (eds), *Advances in Neural Information Processing Systems* 7. Cambridge, MA: MIT Press, 617–624.

- Gibson, E. J. (1969). *Principles of Perceptual Learning and Development*. East Norwalk, CT: Appleton-Century-Crofts.
- Gibson, J. J. (1979). *The Ecological Theory of Perception*. Boston: Houghton and Mifflin.
- Gillebert, C. R., Op de Beeck, H. P., Panis, S., and Wagemans, J. (2008). Subordinate categorization enhances the neural selectivity in human object-selective cortex for fine shape discriminations, *Journal of Cognitive Neuroscience*, *21*, 1054–1064.
- Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*, 178–200.
- Goldstone, R. L. (2003). Learning to perceive while perceiving to learn. In R. Kimchi, M. Behrmann, and C. Olson (eds), *Perceptual Organization in Vision: Behavioral and Neural Perspectives* (pp. 233–278). Mahwah, NJ: Lawrence Erlbaum Associates.
- Goldstone, R. and Barsalou, L. W. (1998). Reuniting cognition and perception: The perceptual bases of rules and similarity. *Cognition*, *65*, 231–262.
- Goldstone, R. L. and Hendrickson, A. T. (2010). Categorical Perception. *Interdisciplinary Reviews: Cognitive Science*, *1*, 65–78.
- Goldstone, R. L. and Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General*, *130*, 116–139.
- Goldstone, R. L., Landy, D. H. and Son, J. Y. (2010). The education of perception. *Topics in Cognitive Science*, *2*, 265–284.
- Huang, X., Lu, H., Zhou, Y., and Liu, Z. (2011). General and specific perceptual learning in radial speed discrimination, *Journal of Vision*, *11*, 1–11.
- Kourtzi, Z. and DiCarlo, J. J. (2006). Learning and neural plasticity in visual object recognition. *Current Opinion in Neurobiology*, *16*, 152–158.
- Landy, D. and Goldstone, R. L. (2007). How abstract is symbolic thought? *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *33*, 720–733.
- Lane, H. (1964). Motor theory of speech perception: A critical review. *Psychological Review*, *72*, 275–309.
- Law, C. T. and Gold, J. I. (2009). A reinforcement learning rule can account for both the associative and perceptual learning on a visual discrimination task, *Nature Neuroscience*, *12*, 655–663.
- Levin, D. T. (2000). Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology: General*, *129*, 559–574.
- Li, W., Pièch, V., and Gilbert, C. D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nature Neuroscience*, *7*, 651–657.
- Lieberman, A. M., Harris, K. S., Hoffman, H. S., & Griffith, B. C., (1957). The discrimination of speech sounds within and across phoneme boundaries. *Journal of experimental psychology*, *54*(5), 358–368.
- MacPherson, F. (2012). Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research*, *84*, 24–62.
- Maurer, D. and Mondloch, C. (2004). Neonatal synesthesia: A re-evaluation. In L. Robertson and N. Sagiv (eds), *Attention on Synesthesia: Cognition, Development and Neuroscience* (pp. 193–213). Oxford: Oxford University Press.
- Moscovici, S. and Personnaz, B. (1991). Studies in social influence: VI. Is Lenin orange or red? Imagery and social

influence. *European Journal of Social Psychology*, 21, 101–118.

Needham, A., Dueker, G.L., and Lockhead, G. (2005). Infants' formation and use of categories to segregate objects. *Cognition*, 94, 215–240.

Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115, 39–57.

Olshausen, B. A. and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607–609.

Peterson, M. A. and Gibson, B. S. (1994). Must figure–ground organization precede object recognition? An assumption in peril. *Psychological Science*, 5, 253–259.

Peterson, M. A. and Lampignano, D. W. (2003). Implicit memory of novel figure–ground displays includes a history of cross-border competition. *Journal of Experimental Psychology: Human Perception and Performance*, 29, 808–822.

Puel, J. L., Bonfils, P. and Pujol, R. (1988). Selective attention modifies the active micromechanical properties of the cochlea. *Brain Research*, 447, 380–383.

Pylyshyn, Z.W. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341–423.

Pylyshyn, Z. (2003). *Seeing and Visualizing: It's not what you Think*. Cambridge, MA: The MIT Press.

Quine, W. V. (1977). Natural kinds. In W. V. Quine, *Ontological Relativity and Other Essays*. New York: Columbia University Press: 114–138.

Quinn, P. C., and Bhatt, R. S. (2006). Are some Gestalt principles deployed more readily than others during early development? The case of lightness versus form similarity. *Journal of Experimental Psychology: Human Perception and Performance*, 32, 1221–1230.

Raftopoulos, A. (2001). Perceptual learning meets philosophy: cognitive penetrability of perception and its philosophical implications. In J. D. Moore and K. Stenning (eds), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 802–808). Mahwah, NJ: Lawrence Erlbaum.

Repp, B. H. (1984). Categorical perception: Issues, methods, findings. *Speech and Language: Advances in basic research and practice*, 10, 243–335.

Rumelhart, D. E. and Zipser, D. (1985). Feature discovery by competitive learning. *Cognitive Science*, 9, 75–112.

Salapatek, P. and Kessen, W. (1973). Prolonged investigation of a plane geometric triangle by the human newborn. *Journal of Experimental Child Psychology*, 15, 22–29.

Schyns, P. G., Goldstone, R. L., and Thibaut, J. P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, 21, 1–54.

Sekuler, A. B., Palmer, S. E., and Flynn, C. (1994). Local and global processes in visual completion. *Psychological Science*, 5, 260–267.

Shiffrin, R. M. and Lightfoot, N. (1997). Perceptual learning of alphanumeric-like characters. In R. L. Goldstone, P. G. Schyns, and D. L. Medin (eds), *The Psychology of Learning and Motivation, Volume 36* (pp. 45–82). San Diego: Academic Press.

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

Smith, L. B. and Kemler, D. G. (1978). Levels of experienced dimensionality in children and adults. *Cognitive Psychology*, 10, 502–532.

- Sowden, P. T., Davies, I. R. L., and Roling, P. (2000). Perceptual learning of the detection of features in X-ray images: A functional role for improvements in adults' visual sensitivity? *Journal of Experimental Psychology: Human Perception and Performance*, 26, 379–390.
- Spector, F. and Maurer, D. (2009). Synesthesia: A new approach to understanding the development of perception. *Developmental Psychology*, 45, 175–189.
- Swoyer, C. (2003). Relativism and the Constructive Aspects of Perception, Supplement to Relativism. In N. Zalta, (ed.), *The Stanford Encyclopedia of Philosophy*. Retrieved April 2011 from <<http://plato.stanford.edu/entries/relativism/supplement1.html>>.
- Tanaka, J. W. and Curran, T. (2001). A neural basis for expert object recognition. *Psychological Science*, 12, 43–47.
- Tanaka, J. W. and Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 46A, 225–245.
- Watanabe, T., Náñez, J. E., and Sasaki, Y. (2001). Perceptual learning without perception. *Nature*, 413, 844–848.
- Weinberger, N. M. (1993). Learning-induced changes of auditory receptive fields. *Current Opinion in Neurobiology*, 3, 570–577.
- Wheeler, D. D. (1970). Processes in word recognition. *Cognitive Psychology*, 1, 59–85.
- Xiao, L.-Q., Zhang, J.-Y., Wang, R., Klein, S. A., Levi, D. M., and Yu, C. (2008). Complete transfer of perceptual learning across retinal locations enabled by double training. *Current Biology*, 18, 1922–1926.
- Yang, T. and Maunsell, J. H. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *Journal of Neuroscience*, 24, 1617–1626.
- Zemel, R., Behrmann, M., Mozer, M. C., and Bavelier, D. (2002). Experience-dependent perceptual grouping and object-based attention. *Journal of Experimental Psychology: Human Perception and Performance*, 28, 202–217.

Notes:

- (1) The 170 in N170 means that the response occurs 170 milliseconds after stimulus onset

Robert L. Goldstone

Robert L. Goldstone, Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN

Lisa A. Byrge

Lisa A. Byrge, Indiana University

