CHAPTER SIXTEEN

Final Commentary

Conceptual Development from Origins to Asymptotes

Robert L. Goldstone and Mark K. Johansen

Scientists studying adult concept learning are typically careful to analyze the entire pattern of responses given across all of the trials of an experiment. Oftentimes the early trials are the most diagnostic because categorization accuracy quickly reaches an asymptote. We take some pride in tackling the hard problem of accounting for adaptive processes that account for category learning, unlike many psychophysicists, who simply throw out the first 1,000 trials because steady-state performance has not yet been reached. However, lest we grow too smug, the chapters of this book provide a great service by reminding us that even though we analyze the very first trial of our experiment, we are still studying conceptual change that occurs almost imperceptibly close to the asymptote. By the time that our 20-year-old subjects come to our laboratories, they have learned the majority of the concepts that they will ever learn and virtually all of their truly foundational concepts. Relatively brief laboratory training suffices to teach students the rule "circle above square" (Bruner, Goodnow, & Austin, 1956), a particular configuration of nine dots (Posner & Keele, 1968), or a new fact such as that grebes are birds, but this rapid learning is only possible because it builds upon a longer and more profound process by which concepts such as *Above* (chap. 3), *bird* (chap. 10), *animal* (chaps. 5, 14), and *animacy* (chaps. 7, 13) are learned.

Those of us who want to develop theories of the learning and representation of adult concepts cannot afford to remain blind to the conceptual development that makes possible adult concept use. This lifelong learning provides us with the fundamental representations that we subsequently combine and tweak. In assessing the contribution of developmental research on concepts and categories to our general understanding of human concepts, we will ask four questions. What are concepts? What is the relation between perception and concepts? What are the constraints on concept learning? What are promising future directions for research on concepts?

What Are Concepts?

A good starting place is Edward Smith's (1989) characterization of a concept as "a mental representation of a class or individual and deals with *what* is being represented and *how* that information is typically used during the categorization" (p. 502). It is common to distinguish between a concept and a category (e.g., Hampton & Dubois, 1993). A concept refers to a mentally possessed idea or notion, whereas a category refers to a set of entities that are grouped together. The concept *dog* is whatever psychological state signifies thoughts of dogs. The category *dog* consists of all the entities in the external world that are appropriately categorized as dogs. In short, concepts are in minds while categories are sets in the external world.

The question of whether concepts determine categories or vice versa is an important, foundational controversy and an extension of the longstanding debate in philosophy over whether concepts correspond to something in the world or are convenient habits of mind. If one assumes the primacy of external categories of entities (the realist and empiricist positions in philosophy), then one will tend to view concept learning as the enterprise of inductively creating mental structures that predict these categories. One extreme version of this view is the exemplar model of concept learning (Estes, 1994; Medin & Schaffer, 1978; Nosofsky, 1984), in which one's internal representation for a concept is nothing more than the set of all of the externally supplied examples of the concept to which one has been exposed. If one assumes the primacy of internal mental concepts (the nominalist and rationalist positions in philosophy), then one tends to view external categories as the end product of applying these internal concepts to observed entities. An extreme version of this approach is to argue that the external world does not inherently consist of rocks, dogs, and tables; these are mental concepts that organize an otherwise unstructured external world.

The authors of this book's chapters generally acknowledge the bidirectional influences between external categories and internal concepts. Some of the important differences of opinion can be expressed in terms of relative focus on one of these influences. Gelman and Koenig (chap. 13), Gopnik and Nazzi (chap. 12), and Mandler (chap. 5) stress the insufficiency of pure inductive learning from perceptually available information for establishing concepts. Mareschal (chap. 4), Rakison (chap. 7), Smith, Colunga, and Yoshida (chap. 11), and Younger (chap. 4) stress the unexpected power of the environmentally available cues that can be used to ground concepts. Still, there appears to be a convergence from all sides on the ideas that tabula rasae cannot form rich and properly constrained concepts and that our abstract concepts are tuned by feedback from perceptual sources. True conceptual learning is not contrary to internal structural constraint. Rather, it is only by having a properly constrained architecture that profound conceptual reorganization becomes possible. This lesson is learned anew every time a first-year graduate student in computer science sets out to hook a camera up to a computer, expose it to years of television input, and have the computer spontaneously organize itself. Flexible concept learning requires sophisticated processes and structures to get off the ground. The search space of possible concepts is incredibly large compared to the set of concepts that humans have found to be functional, and it is easy to forget about all the nonfunctional concepts that humans do not have precisely because we do not have them. Both metaknowledge and domain knowledge, either developmentally or evolutionarily learned, can provide structural constraints. These constraints serve to bias the search in concept space for functional concepts, that is, concepts that neither over nor undergeneralize and are thus maximally useful.

The Functions of Concepts

In assessing which concepts children have, it is useful to bear in mind some important functions of concepts. First, concepts are equivalence classes. In the classical notion of an equivalence class, distinguishable stimuli come to be treated as the same thing once they have been placed in the same category (Sidman, 1994). This kind of equivalence is too strong when it comes to human concepts because even when we place two objects into the same category, we do not treat them as the same thing for all purposes. Still, it is impressive the extent to which perceptually dissimilar things can be treated equivalently, given the appropriate conceptualization. To the biologist armed with a strong *mammal* concept, even whales and dogs may be treated as equivalent in terms of predictions about anatomy and physiology, infant feeding and rearing, and thermoregulation. Once one has formed a concept that treats all skunks as equivalent for some purposes, such as when a child hears several distinct animals labeled with the noun *skunk*, irrelevant variations among skunks can be greatly deemphasized and commonalities emphasized (Keil, 1989; chaps. 9, 13). Several of the chapters make recourse to this notion of stimulus equivalence, arguing that evidence for categorization by young children exists when perceptually discriminable stimuli are nonetheless treated similarly (chaps. 3, 5).

Second, concepts function as building blocks for complex thought. Just as an endless variety of architectural structures can be constructed out of a finite set of building blocks, so concepts act as building blocks for new concepts and an endless variety of complex thoughts. Once a concept has been formed, it can enter into compositions with other concepts. Studying how novel combinations of concepts are produced and comprehended is an active area in adult cognition (Murphy, 1988; Wisniewski, 1997). Similarly, once a concept has been established in childhood, it is subsequently used to build more sophisticated concepts (chaps. 7, 10, 13).

Third, concepts facilitate communication, which among other things allows people to acquire knowledge that enables them to make functional predictions in situations that they have not personally experienced. Communication between people is enormously facilitated if they can count on a set of common concepts (consider "Hey, there's a bear by the beehive"). People's concepts become more consistent and systematic over time in order to unambiguously establish reference for another individual with whom they need to communicate (Garrod & Doherty, 1994; Markman & Makin, 1998). Given the use of concepts for communication, it is not surprising that there are strong ties between concept learning and word learning (chaps. 2, 9–13). Waxman (chap. 9) argues that children learn not only word-to-object mappings but word-to-category mappings, making explicit the important role that concepts play in mediating between objects and words. Gelman and Koenig (chap. 13), Gopnik and Nazzi (chap. 12), and Waxman (chap. 9) show specific ways in which developing concepts not only provide a basis for language but are themselves transformed by language. Concepts provide a foundation for communication, but they do this not by remaining rigid and fixed (like the foundation of a house) but by adapting to the needs of communication (like the foundation supplied by a good pair of sandals) (see also Goldstone & Steyvers, 2001; chap. 11).

Fourth, and most important, concepts facilitate predictive inference by allowing people to generalize experience with some objects to other objects from the same category. Experience with one slobbering dog may lead one to suspect that an unfamiliar dog may have the same proclivity. Oakes and Madole (chap. 6) and Gelman and Koenig (chap. 13) take this inductive generalization function of concepts to be particularly important. Consistent with their focus on this function, the concepts that are most salient and probably most used by humans are exactly those that allow many properties to be inductively predicted.

Categories can be arranged roughly in order of their grounding by perceptual similarity: natural kinds (*dog* and *oak tree*), manmade artifacts (*hammer*, *airplane*, and *chair*), ad hoc categories (*things to take out of a burning house*, and *things that could be stood on to reach a light bulb*), and abstract schemas or metaphors (e.g., *events in which a kind action is repaid with cruelty*, *metaphorical prisons*, and *problems that are solved by breaking a large force into parts that converge on a target*). For the latter categories, members need not have very much in common at all. An unrewarding job and a relationship that cannot be ended may both be metaphorical prisons, but the situations may share little other than this. In contrast to these categories, most natural kinds and many artifacts are characterized by members that share many features (Rosch & Mervis, 1975), and as such they permit many inductive inferences. If we know something belongs to the category *dog*, then we know that it probably has four legs and two eyes, wags its tail as a sign of friendliness, is somebody's pet, pants, barks, is bigger than a breadbox, and so on.

Conversely, the features "has four legs," "tail wags," "barks," and "chases cats" tend to co-occur. So instead of saying "A four-legged, tail-wagging, bark-capable cat chaser crossed the road," little information and corresponding predictive inference potential is likely to be lost by saying "A dog crossed the road." In the same way that MP3 and zip files compress music and data by exploiting repetitions/redundancies in the files, concepts are a way of compressing reality's phenomenal stream by exploiting some of the repetitions that make reality predictable (Bar-Yam, 1999). Interestingly, in the adult category learning literature, Feldman (2000) has recently found a relationship between the difficulty of learning a set of arbitrary categories and the complexity or minimal descriptive length of the shortest rule specifying the two categories. He found that the more compressible the categories, that is, the shorter the description of the rule specifying them, the easier they are to learn. Predictability, learnability and compressibility are closely interrelated. Reality, unlike history, is destined to repeat itself only for those who know the right concepts.

Evidence for Concepts

Research on adult concept learning is dominated by a relatively small number of techniques to uncover conceptual representation. By far the most typical technique is to tell a subject "There are two categories to learn. When you see an item, make a guess as to its categorization. You will receive feedback indicating whether your categorization was correct." Another popular technique is to have subjects list features associated with a verbally presented concept. A third popular technique is to tell subjects that a thing has a certain property and ask them to indicate what is the likelihood that another thing has the property. This final technique has been successfully used with children less than 5 years of age (chaps. 12, 13), but oftentimes developmental psychologists have been forced by the limited attention spans, linguistic capacities, and meta-analytic abilities of their subjects to develop novel techniques for measuring conceptual knowledge. As Oakes and Madole (chap. 6) argue, the requirements of a categorization task influence conceptual representation, and the method of measuring a concept influences what concept a child is inferred to have. From this perspective, it is important to analyze the nature of the methods used to reveal concepts in children.

Developmental techniques, particularly those used with infants below the age of 2, are often indirect measures of concept learning that are related to generalization. Quinn (chap. 3) uses the preferential looking and generalization paradigm. Exemplars of one category (such as "triangle above line") are shown to infants, and the experimenter measures dishabituation to novel exemplars, with increased looking time suggesting that the exemplar is in a different category for the child. While Quinn (chap. 3) finds increased attention for novel rather than familiarized objects, Jusczyk (chap. 2) reports a considerable corpus of evidence from auditory speech processing that requires the assumption that familiarized objects attract a children's attention more than novel objects. These patterns of results are not irreconcilably inconsistent because of many differences in their paradigms. However, the discrepancy does highlight the importance of developing mechanistic process models underlying children's habituation, such as the type developed by Mareschal (chap. 14). Mandler (chap. 5) and Rakison (chap. 7) describe a technique that assumes that objects that are sequentially touched or sorted together are placed in a common category. Mandler and colleagues have used generalized imitation, where the infant is shown an object interacting with another object and encouraged to imitate the interaction with new objects. Waxman (chap. 9) explores generalization by labeling one object and observing what other objects a child selects when the label is repeated. Applying learned label-object associations in the reciprocal direction, Gopnik and Nazzi (chap. 12) use children's likelihood of applying the same label to two objects as a measure that they have been categorized together.

A noteworthy feature of several of these measures is that they do not require that the child has any established categories or concepts at all! Showing that a response generalizes from object A to B but not C does not show that A and B are in the same category, either in the sense that A and B are treated as equivalent, or that children have a preestablished cluster, definition, or characterization that includes both A and B but not C. Generalization is a truly universal phenomenon (Shepard, 1987) that only requires that objects be linked by similarity relations, not by categories. This is particularly true for "single-category" experiments in which a number of objects are familiarized and response generalization to other objects is observed. Such response generalization can occur on the basis of similarity between the familiarized and test objects even if none of the objects are categorized at all. This possibility is consistent with Mareschal's neural network modeling (chap. 14), in which concepts are not explicitly tokened in the network but rather are emergent and implicit results of stimulus processing. It is also consistent with Rakison's notion (chap. 7) that concepts are constructed on-line during categorization rather than having fixed and singular characterizations.

If simple response generalization does not indicate use of concepts or categories, and transfer of a learned label from a known to novel object can be viewed as simply another example of response generalization, then why is it ever necessary to talk about children possessing concepts? One answer is that once words are used to encompass a set of objects, the set seems to have more cohesion for the child. The child becomes interested in knowing what attributes characterize, if not define, the object brought together by a label and may be disposed to look for hidden, abstract, or theory-driven organization principles (Murphy & Medin, 1985; chaps. 5, 12, 13). Originally graded generalization responses give way to more discrete, categorical responses as a child gets more information about the members and attributes associated with a category, reflects on the commonalities possessed by the category members, and works to develop an economical and efficient representation to capture these commonalities.

What Is the Relation between Percepts and Concepts?

One of the primary tensions among this book's chapters, noted by Rakison (chap. 7) is in reconciling the roles of perceptual and nonperceptual information in forming concepts. Several of the chapters react against the intuitively compelling notion that there is a development trend from perceptually based concepts to abstraction-based ones. Several authors argue that even young children's concepts are based on causal, abstract, functional, theory-driven principles (chaps. 5, 7, 12, 13). Taking perhaps the most radical stance along these lines, Mandler (chap. 5) argues against a perceptual-to-conceptual trend because "if anything, development proceeds from the abstract to the concrete, rather than the other way around" (p. 104). Rakison (chap. 7) argues against the trend because he sees a common association-learning process that continues throughout childhood and is responsible for what appears to be increasingly abstract concepts. Likewise, Oakes and Madole (chap. 6) argue that children start off attending to few perceptual dimensions but gradually attend to more as they develop new ways of interacting with their world and then finish the U-shaped trajectory by attending to fewer perceptual dimensions because they use inductively acquired domain knowledge to constrain their attention.

Like these authors, we advocate a revision of the standard perceptual-to-conceptual trend. Our perspective is that perception and abstraction are better viewed as being related by a cycle rather than by a linear continuum. Reminiscent of proposals by several authors, including Rakison (chap. 7), Smith et al. (chap. 11), and Waxman (chap. 9), we will argue that *abstract* construals do not typically supplant perceptual ones but rather that abstract and perceptual construals mutually inform one another.

One of our main reasons for rejecting a simple perceptual-to-conceptual trend is that our developing categories not only are influenced by perception but also influence perception. Intentions, strategies, concepts, and knowledge all affect the perception of similarities among objects. Researchers have found that similarity assessments for a set of objects are affected by learned categorizations of those objects (Goldstone, Lippa, & Shiffrin, 2001; Lassaline, 1996; Livingston, Andrews, & Harnad, 1998). Although similarity ratings are strategic and sophisticated judgments themselves (Goldstone, 1994b), effects of categorization have also been found on tasks that tap more elementary perceptual processes like physical same/different judgments (Goldstone, 1994a) and part detection (Lin & Murphy, 1997). Electrical brain signals that occur as early as 170 milliseconds after the onset of a stimulus are affected by extended training (Fahle & Morgan, 1996) and lifelong expertise (Tanaka & Curran, 2001). Data such as these suggest that perception is adapted to promote the categories or responses required for performing a task, and these adaptations often occur at an early stage of processing (for a review, see Goldstone, 1998).

Our claim is that the development of categories, knowledge, and expertise frequently affects perceptual processing, particularly when these higher-level aspects are frequently important over an extended period of time, as is true of time scales that interest developmental psychologists. This may sound like a controversial claim for top-down processing, given the apparent self-enclosed modularity of perceptual systems (Fodor, 1983; Pylyshyn, 1999), but it is important to distinguish between two kinds of top-down processing. By one account, higher-level processes dynamically, and in a moment-to-moment fashion, affect the processing of lower-level processes. McClelland and Rumelhart's (1981) Interactive Activation Model of word perception is a classic example of this architecture, in which word-level and letter-level processing proceed simultaneously and mutually influence one another *as a word is presented.* There are strong limits on the impact that this kind of top-down influence can have, set by the temporal course of processing. Electrical signals require about 10 milliseconds to propagate across one cortical neuron. Given that visual processing as little as 170 milliseconds after stimulus onset is modulated by knowledge (e.g. Tanaka & Curran, 2001), there are not many cycles of dynamic activation passing possible.

However, in situations where dynamic modulation of a bottom-up signal by a top-down source is implausible, a second kind of top-down effect is possible that does not require a gradually activated knowledge source to excite lower levels. A lower-level perceptual process can simply change its processing in response to consistently occurring, higher-level considerations. Such top-down effects can be accommodated by a strictly feed-forward neural architecture without bidirectional activation passing during the course of stimulus processing. Well-documented cases of this kind of change are apparent in the topographical representations of the primary sensory cortex. For example, monkeys trained to make discriminations between slightly different sound frequencies develop larger cortical representations for the presented frequencies than control monkeys (Recanzone, Schreiner, & Merzenich, 1993). Similarly, monkeys learning to make a tactile discrimination with one hand develop a larger cortical representation for that hand than for the other hand (Recanzone, Merzenich, & Jenkins, 1992). Elbert et al. (1995) measured brain activity in the somatosensory cortex of violinists as their fingers are lightly touched. There was greater activity in the sensory cortex for the left hand than the right hand, consistent with the observation that violinists use their left-hand fingers considerably more than their right-hand fingers. All that is required for these kinds of neural plasticity is that habitually important discriminations become sensitized. Although this is a commonplace form of perceptual learning, it also provides a mechanism by which the perceptual system is tuned to what is functionally important for an organism.

The argument that perceptual similarity is powerful because it can be tuned to an organism's needs is a two-edged sword. Turned around, a critic can argue that the flexibility of perception only exposes its inadequacy as a solid ground for explaining cognitive processes. Certainly, perceptual similarity's explanatory value is attenuated if it is based on exactly those processes that it attempts to explain (similar arguments are presented by Goodman, 1972). However, perceptual processing is slower to change than higher-level conceptual processing, and usually people cannot strategically change their own low-level perceptual processing. Perceptual processing changes, but principally because of the statistical regularities found in the environment (chaps. 2, 4). Transitory conceptions or task-specific needs will not typically modify perceptual systems permanently. However, if a task-dependent categorization is frequently made, or is particularly promising for its organizing power, then it may eventually change the perceptual similarities that are noticed.

The preceding argument for conceptually motivated perceptual change is a species of a more general argument that perception can ground our concepts because it is more sophisticated than it appears at first (neonatal) glance (Jones & Smith, 1993; chaps. 3, 7, 11, 14). Gelman and Koenig's rejoinder (chap. 13) to this argument is that perceptually available features, no matter how sophisticated, are not sufficient to explain how children see beyond costumes, stress the importance of internal structure (even if they have never seen these structures), respect ontological distinctions, and emphasize causal and theoretically motivated features. For example, Gelman and Koenig argue (chap. 13) that the concept *dog* is not characterizable as possessing the perceptually defined feature *leg* because other objects (such as tables) also possess legs, and retorting that it must be the "right kind" of leg is empty in that it drains the perceptual account of its explanatory power. We do not wish to argue against Gelman and Koenig's claim that knowledge and theories influence our categories, but we do suspect that children go through a process of figuring out how to refine and tune their perceptual features. A child might perhaps leave out critical information about movement, texture, and detail in their initial characterization of a dog's leg, making discrimination between a dog's leg and a stuffed dog's leg difficult. However, given the importance of distinguishing real dogs from stuffed dogs and animacy from inanimacy (chap. 7), the originally overgeneral description of leg will tend to be refined until it is a useful perceptual cue.

If perception is frequently motivated by and tuned to its function for higher-level cognition, then perception should not be replaced by abstract reasoning as suggested by a perceptual-to-conceptual developmental trend. Instead, the developmental trend is better described as one from default perception to tuned perception. One phenomenon that illustrates the cooperation between perception and cognition is *interpreted perception*, the process of seeing a thing *as* something (Wittgenstein, 1953). As a guppy is seen as a fish, as a pet, as a vertebrate, or as an interior decoration, the perceptual experience of the guppy changes, and it does so in ways systematically related to the observer's perspective and knowledge. For example, Lin and Murphy (1997) showed that adults in a feature detection task were particularly likely to detect features of a category example that were important as determined by their background knowledge. Hence the primary effect of increased knowledge may not be to reduce the cognitive impact of perception. Rather, it may be to enrich perceptual experience and consequently increase its impact.

What Constrains Concept Formation?

Young children are awash in a sea of features that can be used to form concepts, and this sea of features becomes even richer as their perceptual, motor, and linguistic abilities develop (chap. 6). The need for constraints on concept formation is one of the few theoretical points in the relatively young field of developmental categorization that virtually every chapter in this book assumes and bolsters. Unconstrained feature covariance calculations based on similarity do not provide a reasonable basis for concept formation because there are too many possible feature correlations in an infant's environment (Murphy & Medin, 1985), but there is little consensus in the field on what the dominant constraints are and where they come from. Opinions run from learned selective attention to perceptual features to innate domain knowledge. Nevertheless, there are several dominant themes—selective attention, language, and domain knowledge.

Selective Attention and Bootstrapping

Many of the constraints on concept formation discussed in this book can be construed as selective attention. Younger (chap. 4) argues for selective attention to correlated features, and Jusczyk (chap 2) for selective attention to certain sounds involved with speech. Rakison (chap. 7) emphasizes selective attention to motion and object parts. These last cues are noteworthy because they can potentially be selected early in the course of perceptual processing, yet they have consequences for deeper conceptual analysis, such as the determination of animacy, natural kinds, and object identity. The correlation between a simple perceptual property related to motion and the conceptual distinction between living things and manmade objects means that processes that selectively attend to the former can eventually inform the latter. Hence the original inspiration for the conceptual distinction may be perceptual in nature. The generalization of this phenomenon is "perceptual bootstrapping," according to which relatively sophisticated features and concepts emerge from originally crude and superficial processing.

Perceptual bootstrapping plays a major role in the chapters by Jusczyk, Mareschal, Oakes and Madole, Quinn, Rakison, Smith et al., and Waxman. Jusczyk (chap. 2) describes bootstrapping from phonemic categories to grammatical categories, and

Waxman (chap. 9; see also chap. 11), in turn, describes bootstrapping from grammatical categories to semantic distinctions. Rakison (chap. 7) describes bootstrapping from the appearance of bodily appendages to their function. Mareschal (chap. 14) describes bootstrapping from relatively raw image properties to animal categories. The essential requirement for bootstrapping to occur is that there exists a correlation between properties at different levels of sophistication. Psychologists may eschew confounded variables, but they are indispensable to systems that need to learn to increase their sophistication. For example, Younger (chap. 4) demonstrates infants' sensitivity to conjunctions of features, but conjunctions of features will frequently be correlated with overall similarity. If a simple object possesses the same conjunction of features as another object, the two objects will typically be overall similar to each other. Overall similarity, a property that young children seem to be particularly adept at processing (Smith & Kemler, 1978), can be used as an inroad for developing sensitivity to specific conjunctions of features. Despite psychologists' strivings for hygienic experiments, function, appearance, and meaning are tightly related to each other. Learning often proceeds by first responding intelligently for the wrong, superficial reasons and then gradually dispensing with the scaffolding provided by the superficial cues.

Linguistic Constraints

While Mandler (chap. 5) points out the very real dangers of assuming a one-to-one correspondence between words and concepts, the influence of language is so strong that it can override the influence of perceptual similarity for categorization in young children (chap. 12). From the point of view of researchers such as Waxman (chap. 9) and Oakes and Madole (chap. 6), entities with the same label are in the same category or at least have some properties in common. Labels impact selective attention and maybe even perception by emphasizing commonalities among objects that share them. In addition, a novel label is a sign that a member of a new category is present if a set of features can be found to form the basis for a new category (as discussed in chap. 10).

Instead of having to detect the clusters of covarying properties that are widely believed to form the basis for categories (e.g., Rosch & Mervis, 1975; chaps. 4, 7, 11, 14) through raw computation, language is a way of culturally transmitting previous clusters detected by authorities with linguistic ability (chaps. 10, 13). A label is a sign that useful covariance information is present and should be used to form a concept that other people have found to be functional. The label can be used as a core around which the concept is built (chaps. 7, 12), which can considerably reduce the amount of computation needed to find useful concepts: instead of calculating the correlations between each feature and every one of the other N features in the set of all experienced features, resulting in the need for computing and maintaining N^2 correlations (considering only pairwise correlations), only the correlations between each of the N features and the label need to be computed and maintained. Clearly infants can detect correlated features, as shown by, for example, Younger (chap. 4), and many researchers emphasize correlated features as a basis for concepts (Rosch & Mervis, 1975; chaps. 7, 14). In addition, Mareschal's modeling (chap. 14) shows that concepts have the potential to arise when labels are not explicitly provided. Still, the search for correlations needs to be constrained, and cultural domain knowledge embodied in language can provide powerful constraints.

Domain Knowledge and Metaknowledge: Innate and Learned

Domain knowledge constrains concept formation (chaps. 4, 13) and allows a child to generate concepts that are not solely based on perceptual similarity. Sensitivity to some domain-specific information is arguably innate, as with the categorical perception of certain speech sounds discussed by Jusczyk (chap. 2). Other information is learned, such as the relationship between form and function, as discussed by Oakes and Madole (chap. 6) and Mervis et al. (chap. 10). In a similar vein are the findings reported by Gopnik & Nazzi (chap. 12) on grouping objects based on functional rather than perceptual similarity as the child learns about object functions.

Knowledge can be in the form of broad metaknowledge about concept formation. For example, a child needs to at least implicitly realize that objects can be classified into a hierarchy of categories (chap. 10) and unlearn the mutual exclusivity assumption proposed by Markman (1989) and discussed by Gelman and Koenig (chap. 13). Another such principle might be: generalize widely until feedback indicates overgeneralization and then narrow the concept, possibly by splitting it into several new concepts.

Evolutionary learning and cultural learning significantly bolster individual learning. It is tempting to place the responsibility for certain types of domain knowledge in the magical hands of evolution by natural selection over long spans of time. The strong similarity between individual learning and behavioral evolution has been widely noted and is well summarized by Skinner (1966). This similarity suggests that evolutionary learning still needs to be supplemented with constraints that work for individual concept learning. The set of possible concepts is very large, even on an evolutionary time scale. A comparison of lifelong and evolutionary adaptation also suggests circumstances that foster each type of learning. At a first pass, humans seem to live in the same, reasonably fixed world, suggesting that adaptation across generations would be most effective. Indeed, many general environmental factors, such as color characteristics of sunlight, the position of the horizon, and the change in appearance that an approaching object undergoes, 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. Namely, to a large extent, a person's environment consists of animals, people, and things made by people. Animals and people show considerable variability, and artifacts vary widely across cultures. Evolutionary pressures may have been able to build a perceptual system that is generally adept at processing faces (Bruce, 1998), but they could not have hardwired a neural system that was adept at processing a particular face such as John Kennedy's, for the simple reason that there is too much generational variability among faces. Individual faces show variability from generation to generation, and variability is apparent over only slightly longer intervals for artifacts, words, ecological environments, and animal appearances. Thus, we can be virtually positive that hand 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. At least in part, individuals decide for themselves what objects they will be exposed to. The kinds of concept learning that the majority of the authors of this book discuss cannot be simply relegated to evolutionary adaptation, because considerable flexibility and tailoring of concepts is required. The constraints on lifelong learning described throughout this book thus assume critical importance. These constraints allow, rather than prevent (as might be thought), flexible concept learning that can be achieved in a single lifetime.

Learning Overhypotheses

The argument thus far has been that a learning system must have constraints on hypothesis formation in order to learn concepts in a practical amount of time but that a considerable amount of flexibility is still needed because different people face different worlds and tasks. One exciting possibility raised by several of the chapters is that some of the constraints may themselves be learnable. One way to think about this possibility is in terms of Nelson Goodman's (1954) notion of an overhypothesis, a hypothesis of the form "All As are B" where A and B are generalizations of terms used in any other hypothesis we're interested in (see Shipley, 1993, for a psychological treatment). One might have hypotheses that all dogs have four legs, all storks have two legs, and all worms have no legs. Generalizing over both animals and leg number, one could construct an overhypothesis that "all animals of a particular type have a characteristic number of legs." The power of such a hypothesis is that upon seeing only a single six-legged beetle, one can infer that all beetles have six legs. Research indicates that adults use probabilistic versions of overhypotheses such as these (Heit & Rubenstein, 1994).

Gelman and Koenig (chap. 13) explicitly argue for children's use of overhypotheses and consider such overhypotheses to be evidence for the theory-driven nature of concepts. Overhypotheses are also consistent with Mandler's evidence (chap. 5) that young children typically reason with concepts at a more abstract, superordinate level than basic-level categories. These authors do not argue for overhypotheses that are learned, but this is precisely the direction that Smith et al. (chap. 11), Oakes and Madole (chap. 6), Rakison (chap. 7), and Jusczyk (chap. 2) pursue. Smith et al. (chap. 11) argue for the critical role of word learning in developing higher-order hypotheses that go from token-to-token associations to type-to-type associations. For example, learning "that solidity signals the relevance of shape and that nonsolidity signals the relevance of material—for objects and substances never encountered before and shapes and materials never experienced before" (p. 286) involves forming an overhypothesis. Madole and Cohen (1995, discussed in chap. 7) describe how

14-month-old, but not 18-month-old, children learn part-function correlations that violate real-world events, suggesting that older children *acquire* constraints on the types of correlations that they will learn. Jusczyk (chap. 2) describes the role of early language experience in establishing general hypotheses about how stress patterns inform word boundaries. Children are flexible enough to acquire either the constraints imposed by a stress-timed language like English or a syllable-timed language like Italian, but once they imprint on the systematicities within a language, they are biased to segment speech streams into words according to these acquired biases.

Overhypotheses can greatly increase the power of inductive learning and generalization. One reason why inductive learning seems so hopelessly inefficient is that researchers ignore learning of associations between properties at multiple levels of abstraction. A child seeing a penguin is not just learning that penguins are black and white but is also learning about the relations between coloration, shape, behavior, climate, diet, and so on for birds, animals, and natural kinds. Overhypotheses do not release us from our dependency on constraints. In fact, given the unlimited number of abstract descriptions applicable to an observed event, constraints become particularly important in directing us toward useful levels of abstraction. We need constraints to bias our search for associations between birds and coloration When looking at penguins, rather than associations between objects located within 5 miles of us and the number of vowels in their labels. Still, the possibility that overhypotheses can be learned goes a long way toward severing the traditional connection between constraints and innateness. Not only do constraints permit (rather than limit) learning but inductive learning can also foster the construction of strong but flexible constraints.

Future Directions

With an eye toward the future, we describe what we think are some important avenues for future progress in the field of concept learning throughout development. First, assuming that children's concepts are just underdeveloped, less-functional concepts on their way to becoming more-functional adult concepts ignores the possibility that children's concepts may need to address different functional needs than adult concepts. We anticipate future work on the question "What functionality do children's concepts have that help the child survive into later childhood?"

Second, much recent work in the adult categorization field has focused on how adults represent categories. Partly because it is so much harder to determine what concepts a child has, compared to an adult, a lot of child categorization research has focused on what concepts a child has, rather than how those concepts are represented. Nevertheless, nothing so firmly establishes what concepts a mind has as determining how those concepts are represented.

Third, there has been a trend in the adult categorization literature toward computational-process models of category representation and learning. Such process models have many advantages, including that they force a theory to be specific and help to clarify hidden assumptions. For example, Mareschal (chap. 14) and Smith et al. (chap. 11) have used connectionist models to bolster their arguments. Consideration of a formal process model in the context of almost any theory and set of data is fruitful, if for no other reason than making clear that the relationship between the data and the theory is much looser than was originally apparent. This argument for process models is closely akin to the emphasis by Oakes and Madole (chap. 6) on the importance of studying the process of concept acquisition in young children.

Fourth, as discussed earlier in this chapter, the method used to ask the question "What concepts does a child have?" is likely to strongly affect the answer. For example, the various methodologies differ in the emphasis placed on perceptual similarity, partly as a consequence of how interactive the task is for the child. The methodologies used for children of different ages have been strongly guided by pragmatic constraints such as functional and linguistic ability. Nevertheless, it would be informative to see more systematic comparisons of these methodologies for the same children at the same time, presumably at an age when both methods are pragmatically useful. It is possible that some of the evidence for a perceptual-to-conceptual shift can be clarified in the light of comparing changing expedient methodologies.

Fifth, we believe that much of the progress of research on concepts will be to connect concepts to other concepts, to the perceptual world, and to language. One dissatisfaction with the currently popular concept representation methods—including rules, prototype, sets of examples, and category boundaries—is that one can easily be misled into imagining that one concept is independent of others. For example, one can list the exemplars that are included in the concept *bird*, or describe its central tendency, without making recourse to any other concepts. However, it is likely that all of our concepts are embedded in a network where each concept's meaning depends on other concepts, as well as perceptual processes and linguistic labels. The proper level of analysis may not be individual concepts as many researchers have assumed but systems of concepts. The connections between concepts and perception on the one hand and between concepts and language on the other hand reveal an important dual nature of concepts. Concepts are used both to recognize objects and to ground word meanings. Working out the details of this dual nature will go a long way toward understanding of how human thought can be both perceptual and symbolic.

In this chapter, we have described ways in which the power of inductive concept learning can be increased: by adapting perceptual processing to accommodate concept learning, by taking advantage of perceptual/conceptual correlations to bootstrap abstract properties, by embracing constraints, and by adapting these constraints over time. These principles will be only part of the story for how children and adults learn their concepts. Other important mechanisms were described by the previous chapters, but still more work is needed on novel mechanisms for acquiring rich, interconnected, perceptually grounded, and linguistically meaningful concepts. The field of conceptual development may still be in its infancy, but as the chapters in this book testify, this is the period where the most fundamental progress in conceptual organization can be found.

Acknowledgments

Many useful comments and suggestions were provided by Lisa Oakes and David Rakison. This research was supported by National Institute of Health Grant R01 MH56871 and National Science Foundation Grant 0125287. The first author can be reached by electronic mail at rgoldsto@indiana.edu, and further information about the laboratory can be found at http: //cognitrn.psych.indiana.edu/.

References

Bar-Yam, Y. (1999). *Patterns of meaning*. Boston: New England Complex Systems Institute.

Bruce, V. (1998). *In the eye of the beholder: The science of face perception*. New York: Oxford University Press.

Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. New York: Wiley.

Elbert, T., Pantev, C., Wienbruch, C., Rockstroh, B., & Taub, E. (1995). Increased cortical representation of the fingers of the left hand in string players. *Science*, *270*, 307–307.

Estes, W. K. (1994). *Classification and cognition*. New York: Oxford University Press.

Fahle, M., & Morgan, M. (1996). No transfer of perceptual learning between similar stimuli in the same retinal position. *Current Biology*, *6*, 292–297.

Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, 407, 630–633.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press. Bradford Books.

Garrod, S., & Doherty, G. (1994). Conversation, co-ordination and convention: An empirical investigation of how groups establish linguistic conventions. *Cognition*, *53*, 181–215.

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

Goldstone, R. L. (1994b). The role of similarity in categorization: Providing a groundwork. *Cognition*, *52*, 125–157.

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

Goldstone, R. L., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations through category learning. *Cognition*, *78*, 27–43.

Goldstone, R. L., & Stevyers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General*, *130*, 116–139.

Goodman, N. (1954). *Fact, fiction, and forecast*. London: Athlone Press.

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (Ed.), *Problems and projects* (pp. 341–352). New York: Bobbs-Merrill.

Hampton, J., & Dubois, D. (1993). Psychological models of concepts: Introduction. In I. V. Mechelin, J. Hampton, R. S. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (pp. 11–33). London: Academic Press.

Heit, E., & Rubenstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 411–422.

Jones, S. S., & Smith, L. B. (1993). The place of perception in children's concepts. *Cognitive Development*, *8*, 113–139.

Keil, F. C. (1989). *Concepts, kinds and development*. Cambridge, MA: MIT Press.

Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 754–770.

Lin, E. L., & Murphy, G. L. (1997). Effects of background knowledge on object categorization and part detection. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 1153–1169.

Livingston, K. R., Andrews, J. K., & Harnad, S. (1998). Categorical perception effects induced by category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 732–753.

Madole, K. L., & Cohen, L. B. (1995). The role of parts in infants' attention to from-function correlations. *Developmental Psychology*, *31*, 637–648.

Markman, A. B., & Makin, V. S. (1998). Referential communication and category acquisition. *Journal of Experimental Psychology: General*, *127*, 331–354.

Markman, E. (1989). *Categorization and naming in children*. Cambridge, MA: MIT Press.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, *88*, 375–407.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.

Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, *12*, 529–562.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289–316.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 104–114.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353–363.

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

Recanzone, G. H., Merzenich, M. M., & Jenkins, W. M. (1992). Frequency discrimination training engaging a restricted skin surface results in an emergence of a cutaneous response zone in cortical area 3a. *Journal of Neurophysiology*, *67*, 1057–1070.

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

Rosch, E., & Mervis, C. B. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573–605.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317–1323.

Shipley, E. F. (1993). Categories, hierarchies, and induction. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 30), pp. 265–301. Academic Press.

Sidman, M. (1994). *Equivalence relations and behavior: A research story*. Boston: Authors Cooperative.

Skinner, B. F. (1966). The phylogeny and ontogeny of behavior. *Science*, *153*, 1205–1214.

Smith, E. E. (1989). Concepts and induction. In M. I. Posner (Ed.), *Foundations of cognitive science* (pp. 501–526). Cambridge, MA: MIT Press.

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

Tanaka, J., & Curran, T. (2001). A neural basis for expert object recognition. *Psychological Science*, *12*, 43–47.

Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin and Review*, *4*, 167–183.

Wittgenstein, L. (1953). *Philosophical investigations*. G. E. M. Anscombe (Ed.) New York: Macmillan.