Reuniting perception and conception

Robert L. Goldstone[a,*], Lawrence W. Barsalou[b]

[a]*Psychology Department, Indiana University, Bloomington, IN 47405, USA*
[b]*Psychology Department, Emory University, Atlanta, GA 30322, USA*

Abstract

Work in philosophy and psychology has argued for a dissociation between perceptually-based similarity and higher-level rules in conceptual thought. Although such a dissociation may be justified at times, our goal is to illustrate ways in which conceptual processing is grounded in perception, both for perceptual similarity and abstract rules. We discuss the advantages, power and influences of perceptually-based representations. First, many of the properties associated with amodal symbol systems can be achieved with perceptually-based systems as well (e.g. productivity). Second, relatively raw perceptual representations are powerful because they can implicitly represent properties in an analog fashion. Third, perception naturally provides impressions of overall similarity, exactly the type of similarity useful for establishing many common categories. Fourth, perceptual similarity is not static but becomes tuned over time to conceptual demands. Fifth, the original motivation or basis for sophisticated cognition is often less sophisticated perceptual similarity. Sixth, perceptual simulation occurs even in conceptual tasks that have no explicit perceptual demands. Parallels between perceptual and conceptual processes suggest that many mechanisms typically associated with abstract thought are also present in perception, and that perceptual processes provide useful mechanisms that may be co-opted by abstract thought. © 1998 Elsevier Science B.V.

Keywords: Perception; Conception; Perceptual similarity; Symbol systems; Rule-based systems; Similarity-based systems

1. Introduction

Reflecting on the sophistication of human thought, we can be impressed with how

* Corresponding author. Present address: Psychology Building, Indiana University, Bloomington IN 47405, USA

far we've come, or how we got here. The first perspective emphasizes the distance between starting and final states, whereas the second emphasizes their continuity. In adopting the second perspective, we will describe ways in which conceptual thought is grounded in perceptual similarity. Certainly, many concepts are partially organized around perceptual similarities (Rosch et al., 1976), but we will argue further that perceptual processes guide the construction of abstract concepts even when this direct link may not be so obvious. By adopting a perceptual perspective, we certainly do not mean to deny the use of abstract rules. On the contrary, our position is that abstract conceptual knowledge is indeed central to human cognition, but that it depends on perceptual representations and processes, both in its development and in its active use. Completely modality-free concepts are rarely, if ever, used, even when representing abstract contents. In short, concepts usually stem from perception, and active vestiges of these perceptual origins exist for the vast majority of concepts. Thus, we will argue that *both* similarity and concepts have their roots in perception, and that both rely heavily on perceptual mechanisms for their implementation.

As we shall illustrate throughout this paper, perception's usefulness in grounding concepts comes from several sources. First, perception provides a wealth of information to guide conceptualization. Second, perceptual processes themselves can change as a result of concept development and use. Third, many of the constraints manifested by our perceptual systems are also found in our conceptual systems.

The crux of our argument is that conceptual processing shares important computational resources with perception. For example, we will propose that mechanisms used to represent information in perception perform double duty, also representing information in concepts. Whereas standard theories of concepts assume that different representational systems underlie perception and conception, we assume that a common representational system underlies both, at least to a considerable extent. Thus, mechanisms that represent shape, color and location in perception, also represent shape, color and location in concepts. Analogously, we will argue that processing mechanisms, not just representational mechanisms, are common to perception and conception. For example, we will propose that mechanisms used to scan perceptions and mental images are also used to scan the content of concepts. Various other parallels in processing will be proposed as well. Throughout this paper, when we propose that concepts rely on perception, we will essentially be arguing that concepts share representation and processing mechanisms from perception, rather than using independent mechanisms that work according to fundamentally different principles.

1.1. The allure of the perception/conception distinction

An impressive lineage of theorists has drawn a distinction between perceptual and conceptual systems (for a historical review, see Arnheim, 1969). The Pythagoreans saw a fundamental separation of the 'heaven' of conceptual abstractions and the 'earth' of perceptual experience. Plato believed in 'generic forms' that could not be found by induction across perceptual aspects of specific instances. Parmenides was

perhaps the first philosopher to distinguish reasoning from perceiving, on the basis of perceptual illusions (e.g. a straight stick dipped in water appearing bent) that must be overcome by the powers of rationality (Kirk and Raven, 1962).

The Greek tradition has continued into the present day, with the notion that perceptual similarities must be, and at least sometimes are, cast aside when creating categories. Quine (1977) considers it a sign of an advanced science if its theoretical concepts are not based on perceptual qualities. Evidence suggests that scientific conceptualization incorporates increasingly deep, abstract properties, as opposed to perceptual properties, with increasing expertise (Chi et al., 1981). Similarly, part of the notion of the recent 'theory' theory of concepts is that concepts are not organized around clusters of perceptual properties, but rather around organized systems of knowledge (Murphy and Medin, 1985; Medin, 1989; Murphy and Spalding, 1995). Developmental support for this hypothesis indicates that even young children have inchoate theories about concepts that allow them to disregard perceptual similarities. Children can group animals by their names (Gelman and Markman, 1986; Gelman, 1988), hidden internal structure (Carey, 1985), or genetic heritage (Keil, 1989) in manners that conflict with perceptual similarity. In adults, categorization judgments can be dissociated from similarity judgments such that X is judged more similar to Y than to Z, but is still placed in Z's category rather than Y's (Rips, 1989; Rips and Collins, 1993; Kroska and Goldstone, 1996).

Finally, there have been several recent efforts to separate similarity-based from rule-based systems. Whereas similarity-based systems establish concepts on the basis of perceptual similarities between a concept's instances, rule-based systems form concepts from explicit symbolic expressions. Shanks and St. John (1994) and Sloman (1996) review evidence for such a distinction. Evidence for the use of rules comes from occasions when categorization drastically and suddenly changes when a simple instruction is provided, or when superficial similarities can be completely ignored when they are inconsistent with an explicit instruction. According to Sloman, evidence for the existence of both rule- and similarity-based reasoning exists in the form of simultaneous, conflicting judgments in a task, due to contradictory evidence from the separate systems. One hallmark of rules that divorces them from perceptual information is precisely their generality and universal applicability regardless of domain. In discussing criterial evidence for rule use (as opposed to similarity), Smith et al. (1992) argue that when people use rules, they are as accurate with unfamiliar as with familiar material, and with abstract information as with concrete information. All of these researchers believe that they find evidence for the use of such rules in human reasoning.

In short, a wealth of evidence and intuition suggests that superficial perceptual similarity does not always determine categories. According to many philosophers and psychologists, abstract, rule-based reasoning is often at odds with perceptual data and must be marshaled in order to counter the misleading influences of superficial percepts. Still, returning to the Greeks (Kirk and Raven, 1962), Democritus, speaking on behalf of the senses, chastises abstract reasoning thus: 'wretched mind, do you, who get your evidence from us, still try to overthrow us?'

2. Reuniting perception and conception

Like Democritus, we believe that conceptual structures develop from perceptual processes and continue to bear vestiges of this legacy. At one level, this is an uncontroversial statement, given that many of our concepts are clearly characterized by perceptual properties. Perceptual properties are often good indicators of important, concept-defining properties and our perceptual systems have evolved so as to establish useful concepts (Medin and Ortony, 1989). Objects that belong to psychologically important concepts often have similar shapes (Rosch et al, 1976), and our perceptual systems offer a tremendous amount of data that is probably underestimated by the use of overly sparse experimental materials (Jones and Smith, 1993).

Two approaches toward reuniting perception with conception are the *eliminative view* and the *agnostic view*. According to the eliminative view, perceptual representations constitute all knowledge. Human knowledge contains no non-perceptual representations. According to the agnostic view, human knowledge has major perceptual components and may or may not also contain non-perceptual components. One of us has developed the eliminative view elsewhere (Barsalou, 1993; Barsalou et al., 1993; Barsalou and Prinz, 1997; Prinz and Barsalou, 1997). Thus, we do not pursue this view here, except to summarize it briefly as a boundary case in the realm of possibilities. Instead, we focus on the agnostic view that perceptual representations are central to conceptual knowledge, without making a commitment to whether non-perceptual representations exist as well. In making this argument, we address both similarity and rules, demonstrating that each has important perceptual origins.

2.1. The eliminative view

Prior to the twentieth century, theories of mind typically assumed that human knowledge is inherently perceptual. Not only did the British Empiricists believe this, but so did most other theorists of mind thereafter, including philosophers such as Kant in 1787 (Kant, 1965) and Russell in 1919 (Russell, 1956). In the early twentieth century, ordinary language philosophy and behaviorism both attempted to expunge mental states and mechanisms from theories of mind. As part of their strategy to eliminate mentalism, they frequently criticized theories that relied on mental images. When the cognitive revolution occurred 50 years later, theorists were reluctant to view the cognitive mechanisms that they readopted as inherently perceptual. Although theorists once again became comfortable with cognitive mechanisms, they remained wary of mental images. Furthermore, the proliferation of formal languages and the computer metaphor made possible new ways of thinking about knowledge in non-perceptual formats. Whereas earlier approaches assumed that aspects of perceptual states become stored in memory to form concepts, these new approaches assumed that perceptual states are transduced into amodal symbols. Much like the words of a language, amodal symbols are assumed to bear arbitrary relations to perceptions and to their referents in the world. Attractive properties of this approach included the abilities to form propositional representations of the world produc-

tively, to represent abstract as well as concrete concepts and to implement these languages on computer hardware.

Early critics in the twentieth century often construed the perceptual view in limited and overly simplistic manners, a tradition many modern critics have continued. For example, critics often construe the perceptual approach as containing images that are only conscious, that are only drawn from sensory states and that are only holistic. Rejection of the perceptual view is typically based on this formulation of the position. Actually, many other formulations are possible, some of which were the theories that the British Empiricists actually proposed. For example, Locke's 1690 (Locke, 1959) theory assumed that images in knowledge could come from internal cognitive states, not only from sensations of the external world. He also argued that images could be analytic and productive, not holistic and unproductive. More recently, a wide variety of researchers across the cognitive sciences, especially in cognitive linguistics, have proposed increasingly sophisticated and powerful theories of perceptually-based knowledge.

As an example of these more modern views, consider *perceptual symbol systems* (Barsalou and Prinz, 1997; Prinz and Barsalou, 1997; Prinz, 1997; (for earlier formulations, see Barsalou, 1993; Barsalou et al., 1993)). This theory begins with the assumption that perceptual representations are not necessarily conscious images but are unconscious states of perceptual systems specified neurally. For example, the representation of a chair might be specified as a configuration of neurons active in the visual system rather than as a conscious mental image. These perceptual representations are not necessarily holistic. Instead, a perceptual representation can be a schematic aspect of a perceptual state extracted with selective attention and stored in long-term memory. For example, selective attention might focus on the form of an object, storing only its shape in memory and not its color, texture, position, size and so forth. This schematic extraction process not only operates on sensory states, it also operates on internal mental events, extracting aspects of representational states, cognitive operations, motivational states and emotions. Once these schematic perceptual representations become established in memory, they can function as symbols. They can refer to entities in the world, they can combine productively using combinatoric and recursive mechanisms and they can implement propositional construals of situations. Furthermore, they can represent abstract concepts such as *truth*, *negation* and *disjunction* by capitalizing on perceptual symbols for internal mental events and simulated external events. Note that this very brief review simply serves to provide a sense of the theory's coverage. We do not present it more fully here, because our goal is not to explicate and defend the eliminativist view. This theory is pursued further by Barsalou and Prinz (1997).

One purpose of developing the eliminativist position is to establish an existence proof that a completely perceptual approach is sufficient for establishing a fully functional symbolic system. If this approach can implement reference, productivity, propositions and abstract concepts, it would be comparable with amodal symbol systems in expressive power. As a result, amodal symbols could be eliminated because they are not necessary. Of course, many sorts of evidence must be considered to determine whether eliminativism is justified, but, again, the primary purpose

of this theory is to serve as an existence proof that one *can* develop a fully functional symbolic system that is inherently perceptual.

If one were to push for the eliminativist view, additional sources of evidence could be brought to bear, besides the argument that perceptual symbol systems have sufficient expressive power. First, the amodal view suffers serious problems. These include lack of accounts of how amodal symbols become transduced from perceptual states and, conversely, for how reference from amodal symbols to perceptual states is established (Searle, 1980; Harnad, 1987). Furthermore, there is no direct empirical evidence that conceptual symbols are inherently amodal. Instead, the primary evidence for amodal symbols is indirect, namely, systems having sufficient expressive power can be constructed from them. Finally, the amodal view is too powerful. It can explain virtually any finding post hoc, yet fails to predict many perceptually-based phenomena a priori or provide insight into them. In contrast, the perceptual symbols view does not suffer from these problems. It provides natural accounts of how conceptual symbols are linked to perceptual states and there is considerable evidence that conceptual symbols are perceptual (as we review shortly). Furthermore, the perceptual view is falsifiable, given that it predicts close parallels between conceptual and perceptual processing. Indeed, it makes the strong prediction that the human conceptual system shares representational mechanisms with perception. If this view is correct, then perceptual mechanisms should become engaged when people perform conceptual processing.

This, then, is the flavor of the eliminative position. Clearly, theories of perceptual symbols and the evidence for them remain to be developed considerably in many regards. At this time, this approach primarily attempts to provide an existence proof that, in principle, perception and conception can be united in a way that does not require amodal symbols.

2.2. The agnostic view

A more moderate approach to reuniting perception and conception is to propose that perceptual information plays a major role in conceptual knowledge, which may or may not also include amodal symbols. Our primary purpose is to convince the reader of this and *only* this point. In the remaining sections, we review a large diversity of empirical phenomena that implicate perception in conception. One might wish to follow these phenomena to the stronger eliminativist conclusion, but we will not take that route here. Instead, we will remain content to conclude that conceptual processing at all levels reflects perceptual mechanisms in unexpected ways.

3. The many roles of perception in conception

The distance between percepts and many of our concepts may seem insurmountable. However, several properties of perception prove very useful, in fact irreplaceable, in constructing concepts. In the remaining sections, we demonstrate how

concepts rely on implicit information in perceptual representations, how perceptually-based holistic similarity plays important roles in cognition, how learned perceptual similarities become conceptual biases, how various abstractions (including rules) originate in perception, how perceptual simulation can underlie conceptual processing and how various perceptual mechanisms enter into higher cognition.

3.1. Freeloading with analogical systems

Systems that reflect perceptual similarities in their conceptual structures have a major advantage over those that only incorporate amodal representations – the same advantage that analog systems have over digital systems. Perceptual and analogical representations, because they preserve aspects of the external object in a relatively direct way, can represent certain aspects of the represented object without explicit machinery to do so (Palmer, 1978). For example, to decide that a particular couch belongs to the category, *things that will fit through the front doorway*, a good strategy is to manipulate an analog representation of the couch's shape in reference to an analog representation of the doorway. If shape and relative size are preserved in one's representation, then one can be confident that conclusions drawn from mental manipulations will be applicable to the real-world couch.

The alternative to computing with perceptually-based representations is to reason from symbolic representations that either completely remove perceptual information or start with a symbolic representation. One prominent example of the latter approach is Lenat's CYC project, an attempt to build common sense reasoning into a computer by having 'knowledge engineers' input symbol-level knowledge (Lenat and Guha, 1994). The main problem with this approach is that a tremendous number of facts is needed to represent the same information conveyed efficiently by shape. A picture is indeed worth a thousand symbols, provided that there are processes (such as rotating, scanning and zooming) that take advantage of the picture's analog format.

Representations that preserve physical properties are often more efficient than purely symbolic representations because they do not require external constraints to ensure proper inferences. With amodal systems, symbol manipulation must be controlled through external means such as inference rules. The relation 'next to' is symmetric but not transitive, whereas 'taller than' is transitive but not symmetric. These facts can be represented in an amodal systems so that it does not infer that Rich is next to Jerry upon being informed that Rich is next to Julie and Julie is next to Jerry, but precluding this inference would require devising and applying explicit rules to be used for inferring 'next to' relations. The alternative is to represent the locations and heights of Rich, Julie and Jerry in a spatial medium, whereupon the proximity and height relations can be read off directly (Johnson-Laird, 1983). The logical inferences that are admissible depend on the physical properties of described world, and, accordingly, it makes good sense to preserve these physical properties in the first place rather than try to mimic their effect with specially-designed rules that add constraints back to those stripped by amodal representations. Notably, percep-

tual representations are required to identify the content of these specially-designed rules.

Furthermore, it is surprising how many seemingly abstract properties can be computed by analog devices. For example, many people would assume that if one wished to find the correlation between two variables (e.g. the heights and weights of a sample of college students), it would be necessary to symbolically represent the variables by numbers and use mathematical equations to derive a measure of correlation. Dewdney (1985) describes an analog alternative (shown in Fig. 1). Partially drive in a nail on a wood surface at each point representing an individual height-weight pair. Fit a single, thin rod approximately into place between the nails, and attach a rubber band to each nail and the rod (an incorrect placement of the rod with respect to a nail can be diagnosed by complete slackness of the connecting rubber band). When the rod is released, it quickly falls into an equilibrium position such that the angle of the rod represents the best fitting linear regression of height on weight, and total slackness of the rubber bands represents the correlation between height and weight. The overall slackness of the rubber bands, measured by strumming them or by feeling how much force is necessary to jiggle the rod, can be used to guide decisions about the relatedness between the variables without ever using numeric symbols. As another example, the shortest connecting path between points can be found by pounding nails between slightly separated, clear boards and dipping the boards into a soap solution. When lifted from the solution, the trapped soap film will form a minimal spanning tree that could be used to efficiently lay cables between cities.

Analog representations such as these are often efficient representations, because properties 'ride for free' within the representation without explicitly being computed
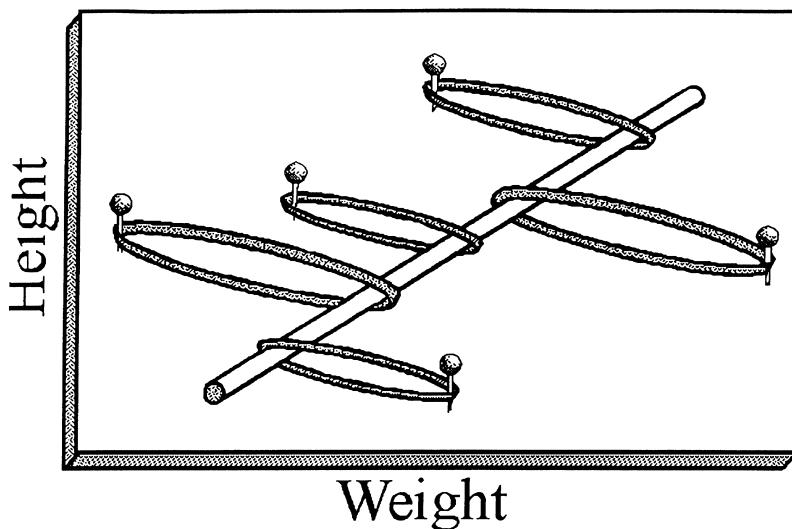

Fig. 1. Although normally computed by amodal mathematical systems, properties such as the correlation between height and weight can also be computed by an analog device consisting of a rod, rubber bands, tacks, and a cork board.

(e.g. correlation, slope and intercept). Analog representations are particularly useful when one does not know what properties will be needed at a later point, or explicitly how to compute the needed properties. As such, ironically enough, analogical representations are often most useful for more complex concepts – those without simple definitions. Although it is unclear what class of object properties are well handled by analog representations, we suspect that many of the properties useful for categorizing objects (e.g. shape, size, pattern of motion, texture, density and curvature) are good candidates. Certainly, other researchers have gotten mileage from analog representations that preserve in a direct form many of the properties and spatial relations present in visual percepts (Shepard, 1984; Finke, 1986).

3.2. The primitive appeal of overall similarity

To the extent that concepts can be characterized thoroughly by simple rules or verbal definitions, the role of perceptual similarity in structuring concepts is weakened. If 'unmarried male human' adequately captures the concept *bachelor*, then sophisticated perceptual representations seem unnecessary.[1] Alternatively, similarity might be said to explain the concept *bachelor*, because members of the *bachelor* category are similar in all being unmarried men, but this hardly salvages similarity as useful for explanation. As Goodman (1972) criticizes similarity: 'when to the statement that two things are similar we add a specification of the property that they have in common, [...] rather than supplementing our initial statement, we render it superfluous' (pp. 444–445). That is, if the similarity of concept members that determines categorization is only with respect to particular properties of commonality, why not just dispense with similarity altogether and discuss the common properties instead?

Our reply to this is to deny the premise. Similarity often involves not only single properties, but integration across many properties. Many of the most important disputes in the field of categorization concern exactly how to integrate across several properties when calculating similarity. Some approaches list features of the two objects, and integrate the overlapping features and the distinctive features to determine overall similarity (Tversky, 1977). Other approaches represent objects as points in a multidimensional space and calculate similarity as an inverse function of the points' distances (Ashby, 1992; Nosofsky, 1992). Still other approaches posit that similarity is proportional to the degree to which the parts of the two compared objects can be aligned with one another (Markman and Gentner, 1993; Medin et al., 1993; Goldstone, 1994a). In all of these approaches, several sources of information (features, dimensions, or parts, respectively), not single criterial properties, are the basis for determining similarity.

Similarity assessments may typically integrate across many properties because it is natural for people to form impressions of overall similarity. In fact, evidence suggests that in many situations, it is easier for people to base similarity and cate-

[1]Although *bachelor* is often presented as one of the clearest cases of an easily definable concept, Lakoff (1986) argues that even here, the definition cannot account for cases such as seven-year old boys or popes being extremely poor examples of *bachelor*, if they are examples at all.

gorization judgments on more, rather than fewer, properties (Kemler, 1983; Goldstone, 1994b). For example, Sekuler and Abrams (1968) report cases in which people are faster to respond that two displays are identical along all their elements than that two displays have a single common element. Nickerson (1972) reviews evidence in favor of a fast, 'sameness detector' that allows people to quickly assess overall similarity between displays before being able to respond to particular dimensions. Individuals whose cognitive judgments are impeded, because they are young (Smith and Kemler, 1978; Kemler, 1983; Smith, 1989), inexperienced (Foard and Kemler Nelson, 1984), hurried (Ward, 1983), or distracted (Smith and Kemler Nelson, 1984), seem to rely on 'holistic' rather than 'analytic' processes. Thus, responding on the basis of overall, undifferentiated similarity may be a more primitive computation than responding on the basis of particular properties.

In a similar vein, Brooks (1978) argued that judging category membership by overall similarity is an often-used strategy, particularly when the category members are rich and multi-dimensional and the category rules are complicated (see also Allen and Brooks, 1991). Determining overall similarity across many properties is efficient relative to determining similarity with respect to a particular property when it is difficult to break an object down into separate features or aspects. Such a state of affairs is likely to occur for many real-world objects. Laboratory stimuli often 'wear their featural compositions on their sleeves,' but natural objects are seldom so obliging (Goldstone and Schyns, 1994; Schyns et al., 1998).

As all of this work illustrates, people have a strong disposition to process overall similarity, and doing so appears to serve a number of important cognitive functions. This is *not* to say that people do not also process similarity more analytically, a topic to which we turn shortly. Nevertheless, people do appear to process similarity holistically on many occasions. We propose that this basic tendency has its origins in perception, and that perception places important constraints on it, thereby mitigating Goodman's problem of unconstrained similarity. The perceptual system is geared toward providing overall similarity because of the importance of parallel comparison in object categorization. Rapid categorization is possible when multiple perceived properties can be matched in parallel to potentially corresponding properties in category knowledge. If processing were serial, processing would probably not be as efficient, and survival might be compromised. Thus, overall similarity may have evolved in perception to aid categorization, having the additional result of producing a comparison mechanism that is used broadly across many other cognitive tasks as well.

Overall similarity seems to be more efficiently processed by perceptual systems than amodal symbolic systems. First, the processes that allow simple features to be registered in parallel without a capacity limit are found widely throughout perceptual systems. For example, the features *red* and *horizontal line* can be detected simultaneously across an entire visual display (Treisman and Gelade, 1980), whereas non-visual semantic features cannot be detected in parallel with unlimited capacity (Mullin and Egeth, 1989). Second, perceptual systems often respond to multiple sources of information by blending them together without individuating the sources. For example, people can combine light from two spatial locations so that

light detection depends on their summed energy (Bearse and Freeman, 1994). Amodal symbols are separately individuated, and thus require explicit combination of symbols in order to produce overall similarity responses. Overall responding can be achieved by analytically combining separated sources of information, or by never differentiating the sources in the first place; perceptual systems do both whereas amodal systems have profound difficulties with the second.

Overall similarity is computed efficiently by perceptual systems, but is useful for concepts more generally. Categories that permit many inductive inferences typically have members that share many features (Rosch and Mervis, 1975). In contrast to the members of metaphor-based categories (such as 'situations that are 'time bombs' waiting to go off' (Ortony, 1979; Glucksberg and Keysar, 1990)) or ad-hoc categories ('things to take from a burning house' (Barsalou, 1991)), our standard taxonomic categories, such as *chair*, *trout*, *bus*, *apple*, *saw* and *guitar*, are structured by overall similarity across many attributes. They are characterized by their effective compromise between within-category similarity and between-category dissimilarity (Rosch, 1975). If we know that something belongs to the category *bird*, then we know that it probably has two legs and two eyes, nests, flies, is smaller than a desk and so on. These inductive inferences are admissible because of the high overall similarity between members of the *bird* category. If we want to rapidly learn and deploy these taxonomic categories, it behooves us to pay attention to overall similarity.

3.3. Learned perceptual similarities

One reason why perceptual similarity is more powerful than might be thought is that it is not inflexible and insensitive to contextual factors. Although similarity certainly affects categorization, there is also an influence, albeit attenuated, in the reciprocal direction. Lassaline (1996) reports that judgments of induction across categories (e.g. if horses have property *X*, how likely is it that cows do?) have an influence on subsequent similarity assessments involving the same items. Kelly and Keil (1987) find that exposure to metaphors can even influence similarity judgments to different materials. For example, subjects who received the metaphor 'The New Yorker is the quiche of newspapers and magazines' gave higher similarity ratings to food-periodical pairs that had similar values on a *tastefulness* dimension (e.g. *steak* and *Sports Illustrated*) than did subjects who were not given these metaphors. In short, impressions of similarity are educated by the more sophisticated tasks that use them.

In the above examples, one could charge that the similarity assessments are influenced by high-level tasks simply because they are quite sophisticated judgments themselves. While similarity ratings do certainly seem sophisticated and cognitively penetrable, effects of categorization have also been found on tasks that tap more perceptually-based similarities. Goldstone (1994c) first trained subjects on one of several categorization conditions in which one physical dimension was relevant and another was irrelevant. Subjects were then transferred to same/ different judgments (e.g. 'are these two squares physically identical?'). Ability to

discriminate between squares in the same/different judgment task, measured by signal detection theory's d'(sensitivity) was greater when the squares varied along dimensions that were relevant during categorization training. In one case, experience with categorizing objects actually decreased people's ability to spot subtle perceptual differences between the objects, if the objects belonged to the same category. Similarly, work in categorical perception indicates that discriminations involving pairs of stimuli that straddle category boundaries are more easily made than are discriminations involving stimuli that fall within the same category, equating for physical dissimilarity between the pairs (Harnad, 1987).

Research in perceptual learning indicates influences of tasks on perceptual systems that are surprisingly early in the information processing sequence. For example, practice in discriminating small motions in different directions significantly alters electrical brain potentials that occur within 100 ms of the stimulus onset (Fahle and Morgan, 1996). These electrical changes are centered over the primary visual cortex, suggesting plasticity in early visual processing. Karni and Sagi (1991) similarly find that the primary visual cortex exhibits adaptation in simple discrimination tasks. In general, similarities are adapted to promote the categories or responses required for performing a task, and these adaptations often occur at an early stage of processing.

The argument that perceptual similarity is powerful because it can be tuned to an organism's needs is a two-edged sword. Turned around, similarity's critic can argue, 'the flexibility of similarity only exposes its inadequacy as a solid ground for explaining cognitive processes'. Certainly, 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; Shanon, 1988; Rips, 1989). However, our position is that perceptual processing is slower to change than higher-level conceptual processing. New conceptualizations, such as categorizing a chair as a device suitable for reaching a light bulb, when the current tasks demand this (Barsalou, 1991), are not immediately transformed into modifications of the perceptual system. Perceptual learning is typically a prolonged process requiring thousands of trials of practice (Shiffrin and Lightfoot, 1997). 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 perceptual similarities that are noticed. Experts eventually come to see the objects in their domain of expertise in a different way than novices (Biederman and Shiffrar, 1987; Burns and Shepp, 1988), and perceptual development in children may involve spotting new perceptual commonalities (Schyns et al., 1998).

In general, then, perceptual similarity may change, but over a small range with a relatively protracted time course. As such, it can serve as a point of departure for highly context-dependent cognitive process such as goal-driven categorization, metaphorical comparison and analogical reasoning. Since perception is flexibly tuned, the departure point itself moves, making often-traveled points more accessible (Goldstone, 1995a). Similarities that were once effortfully constructed, become second nature to the organism. In this manner, perceptual similarity can

provide a useful starting-off point for specialized cognitive processes – useful because it has been tuned, although perhaps slowly, to the tasks that use it.

The distinction between slowly-changing perceptual processes and more labile conceptual processes is closely related to the associative/rule distinction (Smith and Sloman, 1994; Sloman, 1996). We suspect that associative mechanisms are, in general, a large class of relatively automatic processes, and that rule mechanisms are, in general, a large class of relatively controlled processes (Shiffrin and Schneider, 1977). Associative mechanisms tend to be those that process bursts of features made available automatically, either through parallel processing in perception or automatic activation in long-term memory. Rule mechanisms tend to be those that process individual features selectively through serial processing in working memory, using limited attentional resources. Whereas associative mechanisms require much practice to develop, rules can be constructed in an ad hoc manner to produce large immediate changes in performance.

3.4. Abstractions from perception

Until now, we have concentrated on the usefulness of perceptual information in conceptual representations. However, we underestimate the importance of perceptual information in concepts if we restrict ourselves to situations in which it is represented directly. In this section, we consider situations in which perception less directly, but not necessarily less strongly, 'jump starts' concepts, by motivating, informing and providing procedures for abstract thoughts.

If we focus on the current state of an abstract concept, we risk ignoring the simpler, less sophisticated starting point that was necessary for its development. Developmental evidence provides several examples of abstractions evolving out of what appear to have been perceptually-based concepts. For example, infants appear biased to treat parts of a display that move together as belonging to the same object (Spelke, 1990). Once spatially separated parts of a display are joined together in the same object because of their common motion, other object properties, such as edges, smoothly varying contours and uniform coloration, can be detected. The detection of motion is also instrumental in acquiring the distinction between living things and artifacts (Keil, 1989; Gelman, 1990; Mandler, 1992). Not all living things move, and not all artifacts are static, but Keil, Gelman and Mandler suggest that the original inspiration for this latter distinction may be based on movement patterns. In fact, children often treat human-made objects with irregular patterns of motion as being alive. This distinction has a significant impact on concepts, determining how extensively we infer properties from one member of a category to others, and whether we believe the categories are organized around essences. This elevated role in inductive reasoning belies the 'lowly' motion-based origin of the natural kind/artifact distinction.

We do not necessarily disagree with theorists who argue that biological categories are organized around 'theories' that involve genetic heritage, internal structure, birth, death and reproduction (Carey, 1985; Keil, 1989; Rips, 1989). For example, Carey's finding that children and adults are more likely to extend an unfamiliar

property from humans to worms than from humans to toy monkeys may derive from knowledge concerning internal organs. However, we suggest that, in many cases, the original inspiration for constructing the theory comes from relatively simple, perceptual cues such as motion, the presence of internal fluids and so forth. Moreover, perceptual representations of events may underlie many abstract concepts (Barsalou, 1993; Barsalou and Prinz, 1997). In the spirit of the earlier section (Section 3.1) perceptual representations of reproductive events, such as mating and giving birth, may provide large amounts of implicit information that underlie biological theories. Without perceptual knowledge of these events, we suspect that people would have little understanding of biological concepts.

The evolution of abstract concepts from perceptual roots has other applications. Pure mathematical concepts are often originally inspired by perceptual evidence. In the development of number concepts, there is strong evidence that children rely on perceptual representations. (Stigler, 1984; Huttenlocher et al., 1994). Similarly, mathematicians frequently report first creating a visual 'proof' of a theorem, and only subsequently derive the symbolic, and publishable, version, written out in theorems and lemmas (Hadamard, 1949; Barwise and Etchemendy, 1990, 1991; Anderson, 1997). In problem solving, one of the most effective ways of deriving an abstract schema, such as 'overcoming an object by converging weak intensity forces from several pathways onto the object' is to use several concrete examples (Gick and Holyoak, 1983). Categorizations that are based on a complex abstract rule may initially be solved by using perceptual similarity between items to be categorized and known category members (Brooks, 1978; Allen and Brooks, 1991).

Another body of evidence indicates that experimentally-noticed perceptual similarities alter more abstract processes. Requiring subjects to perform similarity judgments on pairs of scenes makes subjects more likely to treat the scenes in an abstract manner subsequently (Markman and Gentner, 1993). For example, early similarity judgments promote responses based on common roles (e.g. 'these two things correspond to each other because they are both donors') rather than on superficial attributes (e.g. 'these go together because they are both men'). As another example, when people are shown an object to be categorized, they are often reminded of a superficially similar object. Once reminded, they try to come up with an abstract description for the category that encompasses both objects (Ross et al., 1990).

In sum, even when concepts eventually come to be characterized by abstractions, these abstractions owe their existence to perceptual similarity. Two consequences follow. First, conceptual end states do not imply an absence of perceptual origins. Even if the end-state of a concept were free of perceptual information, perceptual processing may have been required to build it. Second, perceptually-inspired abstractions can provide a mechanism for developing abstractions not currently within the abstract system's powers. New expressive capacities arise when abstractive processes create new uses and descriptions for concepts that have been established perceptually.

To borrow an example from evolution, where we have good reasons to think that genuinely new structures and functions arise, mammalian ear bones probably evolved from jaw bones (Gould, 1993). The sound-transmitting function of the

jaw/ear bones was only selected for once the biting function of the jaw had already established the jaw bone's basic shape. Likewise, the starting shape of our concepts may be perceptually-specified initially, but can be transformed in quite different directions once developed. Later structures (ear bones and abstractions) depend on and grow out of earlier structures (jaw bones and perceptually-based concepts) for their very existence even as they acquire radically different functions.

In Section 3.6 we suggest several examples of conceptual mechanisms (high-lighting, structuring, productivity) that may have evolved from perceptual mechanisms (e.g. selectivity, binding, juxtaposition). Although Fodor (1975) has argued that it impossible for a symbolic representation system to learn concepts that have expressive powers not already present in the system, we propose that perceptual mechanisms have lead the conceptual system to entertain hypotheses that it would not otherwise have entertained. Our view stresses the evolution of concepts from perception, through processes that can eventually achieve abstract end states.

3.5. *Perceptual simulation in conceptual tasks*

Standard conceptual tasks that lack pictorial materials offer a means of assessing how far perception extends into conceptual processing. In the feature-listing task, for example, subjects receive the word for a concept and verbally list features typically true of its instances (e.g. list features typically true of 'watermelons'). Similarly, in the property-verification task, subjects receive the word for a concept and verify whether a second word specifies a property true of the concept (e.g. for 'watermelon' is 'seeds' a property?). In neither task do subjects receive pictures, nor are they asked to use imagery. Instead, they receive only linguistic materials with neutral instructions.

Perceptual mechanisms play no role in standard theories of these tasks. Instead, these theories assume that subjects access feature lists, frames and semantic nets that only contain amodal symbols. Recent evidence, however, strongly implicates perceptual mechanisms (for a review, see Barsalou et al., 1997). Rather than accessing amodal representations, subjects appear to simulate referents of the concepts perceptually and then scan these simulations to produce the required information.

For example, Wu's (1995) dissertation found evidence of perceptual simulation in the feature listing task. Subjects produced features for nouns (e.g. 'watermelon') and related noun phrases (e.g. 'half watermelon'). Two sources of evidence indicated that subjects simulated the referents of these nouns and noun phrases in order to list features. First, subjects who received neutral instructions produced essentially the same features as subjects asked to construct and describe images, suggesting that the neutral subjects adopted images spontaneously. Second, the visibility of features in real-world referents (not encountered in the experiment) predicted their likelihood of being produced. For example, the feature *seeds* is occluded in the perception of a whole watermelon but is visible in the perception of a half watermelon. Across four experiments, noun phrases, for referents with non-occluded internal properties (e.g. 'half watermelon') produced much higher rates of internal feature listing than nouns for referents with occluded internal properties (e.g. 'watermelon'). These results

strongly suggest that subjects simulated referents of the concepts perceptually in order to produce features.

Perhaps subjects in Wu's experiments performed perceptual simulation because feature listing is a deliberate, recall-oriented task that allows time for such simulations. If so, then we shouldn't observe such effects in faster, recognition-oriented tasks that produce reaction times under 1 s. To explore this issue, Solomon's (Solomon, 1997) dissertation explored the role of perceptual simulation in the property verification task (for a preliminary report, see Olseth and Barsalou, 1995).[2] Although Kosslyn (1976, 1980) found no evidence of perceptual simulation when neutral subjects performed property verification, the easy false trials in his experiments made such simulation unnecessary. Like Kosslyn, Solomon found that when the false trials were easy (e.g. does a *crab* have a *brick*?), subjects adopted a linguistic associative strategy that bypassed conceptual knowledge and perceptual simulation. However, when the false trials were difficult (e.g. does a *crab* have a *fin*?), subjects could not use the linguistic associative strategy and had to use conceptual knowledge instead. Under these conditions, perceptual factors, such as property size, provide the best prediction of reaction time. Thus, when subjects were forced to perform conceptual processing, they resorted to perceptual simulation to find and verify property information about concepts.

Further evidence for this conclusion comes from the second half of Solomon's 1997 dissertation. In these experiments, the perceptual similarity of parts was manipulated across trials. Whereas one subject might first verify the concept-property pair, *PONY-mane*, and later the pair, *HORSE-mane*, another subject might first verify *LION-mane* and later *HORSE-mane*. If subjects perceptually simulate the concepts to verify the parts, then they should be faster for the *PONY-HORSE* sequence than for the *LION-HORSE* sequence, because pony manes are more similar to horse manes than are lion manes. When subjects process the *HORSE-mane* pair, they are reminded of the earlier pair involving *mane*, which either facilitates or inhibits processing. Across three experiments, Solomon observed this result, suggesting that subjects perceptually simulated the concepts to perform verifications. One could argue that this effect results from *HORSE* being more similar to *PONY* in general than to *LION*. To assess this possibility, all experiments included materials in which the part is equally similar to all three concepts. For example, the part, *back*, is roughly the same across *PONY*, *HORSE*, and *LION*. For these materials, there is no difference between sequences such as *PONY-HORSE* and *LION-HORSE*, indicating that part similarity, not concept similarity, is the important factor.

Together, the results from these three projects indicate that perceptual simulation is central to conceptual processing. Even when subjects receive no pictorial materials and are not asked to use imagery, they nevertheless perform perceptual simulation spontaneously.

3.6. Parallels between perceptual and conceptual processing

Contrary to the Greek philosophers' polarized dichotomy between perception and

---

[2]Karen Olseth is now Karen Olseth Solomon.

cognition, we have seen that there is good reason to believe that cognitive processes borrow from perceptual ones. Cognitive economy and evolutionary considerations (large frontal lobes being a relatively recent evolutionary advance) encourage the co-option (borrowing) of perceptual processes for symbolic cognition. Thus, although some symbolic reasoning operations appear to have no perceptual equivalent (e.g. deduction and modus tollens; but see Johnson-Laird (1983) for concrete models of these abstract operations, and Rips (1986) for a reconsideration), many suggestive parallels can nevertheless be drawn elsewhere. The points to be taken from these parallels are: (a) that properties typically associated with abstract cognition are often present in perceptual systems, (b) perceptual systems have mechanisms that are useful for more abstract cognition and provide new insights into how higher-order cognition may operate and (c) patterns of correlations between perceptual and conceptual processes suggest that they share common mechanisms. While the evidence for (c) is admittedly correlational rather than causal, additional considerations with respect to the evolution and development of mental abilities suggest that these correlations may often be due to conceptual processes borrowing from perceptual ones. Table 1 lists several parallels, which we consider next.

3.6.1. Selectivity

A hallmark of abstract, rule-like, cognition is that it emphasizes certain properties over others. To apply the rule 'an island is any piece of land completely surrounded by water' to a particular plot of land, one must emphasize this *criterial* attribute and ignore characteristic island features such as *tropical* and *sandy* (Keil, 1989). This selective highlighting of important attributes has a clear parallel in the considerable body of work on selective attention (for a review, see Johnston and Dark, 1986). Many properties of perceptual attention make it a promising candidate for subserving situations where more abstract highlighting of properties is needed, and indeed recent theories of knowledge and language have incorporated it centrally (e.g. Talmy, 1983; Langacker, 1986; Mandler, 1992; Barsalou, 1993). First and most basically, perceptual selection of relevant information is highly effective. Researchers have found that when people are instructed to respond to one of two overlapping shapes, there is very little processing of the irrelevant shape (Garner, 1974, 1978; Rock and Gutman, 1981; also Melara and Marks, 1990), and very little performance decrement compared with when just one shape is shown (Neisser and Becklen, 1975). Second, attention can be directed to particular stimulus properties (Treisman

Table 1
Parallels between cognitive and perceptual processes.

Cognition	Perception
Highlighting	Selective attention
Abstraction	Blurring, filtering
Structuring	Binding
Differentiation	Dimensionalization
Analogy	Cross-modal matching
Productivity	Image juxtaposition

and Gelade, 1980), and properties automatically capture attention if they have been important during prolonged training (Shiffrin and Schneider, 1977; Logan and Etherton, 1994). This latter capacity could underlie people's ability to learn new criterial, abstract definitions. Third, attention is not only directed by simple stimulus properties, but also by semantic coherence and context (Triesman, 1960). If perceptual selection processes were not capable of being driven by higher-level properties, then it would have had limited application to more strategic cognition.

From this description of perceptual selection, it is clear that, at a minimum, we cannot use selectivity as evidence for symbolic, rule-like cognition. Perception also benefits from a sharp, efficacious form of selectivity.[3] Our critic then continues, 'just because both abstract reasoning and perception have mechanisms for highlighting relevant properties does not mean that the abstract ability derives from the perceptual ability'. Although far from definitive, evidence exists for the co-option of perceptual selective attention by cognitive selective attention. One line of evidence comes from an examination of individual differences and mental disorders. For example, schizophrenic patients have characteristic attentional and cognitive deficits that parallel each other in interesting ways. Cognitive correlates of schizophrenia include abnormal word associations, problems with developing coherent discourses and difficulties with abstract thought that stem from intrusions of superficial information (Schwartz, 1982). Perceptual correlates of schizophrenia include problems with allocating attention to relevant stimulus attributes, driving attention by informative cues and inhibiting irrelevant attributes (Liotti et al., 1993).

The parallel between these high- and low-level deficits is that both involve problems with selective attention. The pattern of correlations is well explained if the same selective attention processes are at work for surprisingly different levels of processing, among tasks that many have claimed to be handled by special modules (e.g. for language and visual information). Although the evidence above is principally correlation, there is further suggestive evidence that perceptually-based selective attention is borrowed for more conceptual selective attention, rather than vice versa. Schreiber et al. (1992) found that schizophrenic patients show impairments on a visual selection task before cognitive impairments arise, even when the tasks are roughly equated for their sensitivity at diagnosing abnormalities. The implication of this result is that if perceptually-based selective attention processes go awry, then more general cognitive impairments of selective attention may follow.

The perceptual and cognitive deficits of schizophrenics seem to be attributed to attentional processes that are insufficiently selective (Beech et al., 1989). That is, schizophrenics have particular difficulty inhibiting inappropriate thoughts and irrelevant perceptions. Conversely, many of the perceptual and cognitive symptoms of childhood autism, including minimizing sensory stimulation, focused attention on a single environmental cue, abnormally narrow generalizations from training and lack of productive language, may be traced to an overly selective attentional process

---

[3]This is not inconsistent with the earlier argument for a primitive sense of overall similarity. We seem to have both effective routines for processing overall similarity and routines for selectively overruling this primitive process, consistent with the automatic versus controlled distinction that we applied earlier to the associative-rule distinction.

(Lovass et al., 1979). A narrowing of attention characterizes both autistic perceptual behavior and autistic use of words and concepts. Sensations and semantic meanings outside of a tightly restricted range are ignored.

Another example of surprising correlations between superficially dissimilar perceptual and conceptual tasks is found in individual differences in 'field (in)dependence' (Wapner and Demick, 1991). Field dependent individuals have difficulty selectively attending to one component of a stimulus while ignoring its context. In the canonical test for field dependence, people locate hidden figures that are embedded in a scene. Most importantly, proficient performance on this perceptual task is correlated with conceptual tasks, such as making fine semantic categorizations, drawing sharp boundaries between concepts, verbal coding strategies and constrained analogical reasoning. In reviewing the evidence for field (in)dependent personalities, Miller (1987) concludes that there are selective attention (and its converse, holistic integration) operations that are shared by a wide array of perceptual and conceptual processes.

Selective attention to important stimulus aspects may emerge from one of two processes – a process that focuses on important, criterial, or goal-based aspects, or a process that actively suppresses irrelevant aspects. While many theories of abstract reasoning focus on the former, recent perceptually-based work has found a strong presence of the latter. For example, research on 'negative priming' (Neill, 1977; Tipper, 1992) has shown that people are slower to respond to a target if it was a distractor on previous trials. Again, non-perceptual equivalents are available. Processes exist to inhibit irrelevant memories (Anderson and Spellman, 1995) and words (Gernsbacher and Faust, 1991), depending on how much they compete with other, potentially more appropriate items.

In short, there appear to be strong correlations between classes of cognitive and perceptual behavior that stem from shared processes of selective attention to relevant properties, and selective inhibition of irrelevant properties. If this speculation is correct, then we might expect interference between cognitive and perceptual tasks that both make demands on the same selective attention process. Focusing on a target location or property (as with Flanker and Stroop tasks) should interfere with focusing on criterial properties in a categorization task. Regardless of the empirical outcome, given the functional similarities between perceptual and conceptual processes for highlighting relevant information, it is plausible that some of the early (phylogenetically and developmentally) processes for allocating perceptual attention would be co-opted for later processes.

3.6.2. Blurring and filtering

Turning to some of the other parallels in Table 1, it is somewhat surprising to note a link between abstraction, often considered the epitome of cognition that has transcended perception, and the 'lowly' processes of blurring and filtering. By blurring, we are referring to any process that removes detailed information from further processing. To abstract is to distill the essence from its superficial trappings. The conventional way to do this is by developing a 'schema' that is tuned to the essence (Gick and Holyoak, 1983). Another way is to blur over the irrelevant

aspects. Blurring has the advantage that it can operate even when the essential schema cannot be formulated. One does not need to know what makes something a dog in order to categorize the neighbor's poodle as a dog, as long as one knows that a beagle is a dog, and is able to ignore (blur over) the differences between poodles and beagles. Visual agnosias provide support for the notion of categorizing by blurring. Often brain damage in visual processing areas results in preserved ability to make coarse categorizations (e.g. 'I know this is some kind of animal'), despite an inability to use narrower categories (Farah, 1990). An agnosic patient may be able to perceive that a photograph depicts some sort of animal without recognizing the animal to be a cat. Similarly, 18-month old children show sensitivity to broader categories such as *vehicle* and *animal* without showing sensitivity to narrower categories like *dog* and *car* (Mandler et al., 1991). Thus, people who have not yet learned how to, or have lost their ability to, categorize at relatively narrow levels, may still be able to make broad categorizations. This suggests that broad categorizations can be obtained simply by not registering within-category differences.

Blurring is a particularly appropriate technique when the superficial aspects are details, and the global structure is correlated with important abstractions. Furthermore, *strategic* blurring may cause only particular features to be ignored, thereby permitting features known to be irrelevant to exert little influence on behavior, even if the relevant features have not been identified. Strategic blurring is probably instantiated by attentional, rather than visual, means, given that the attentional system shows flexibility in allocating intermediate degrees of attention to stimulus features. At times, attention given to a stimulus feature nearly optimally matches the feature's diagnosticity for a task (Shaw, 1982; Nosofsky, 1986). Thus, ignoring is not necessarily all-or-none. Whereas amodal abstractions must either represent or not represent a property, an advantage of incorporating perceptual attention in high-level cognition is that it provides a mechanism for partially representing a property.[4]

3.6.3. Structure and binding

The binding of arguments to values is often proposed as a structure that distinguishes high-level cognition from its lower-level counterpart. The proposition *Loves*(*John*, *Mary*) means something different than *Loves*(*Mary*, *John*), much to John's dismay. The predicate *Loves* takes arguments that are ordered by their roles, with the 'agent' role bound to John, and the 'patient' role bound to Mary. This binding of objects to roles establishes a structure in propositions that goes beyond the representational capacities of 'flat' representations such as feature lists (Barsalou, 1992, 1993; Barsalou and Hale, 1993). Although structured representations are necessary for orderly thought, abstract thought is not alone in this regard.

---

[4]The typical solution to representing partially diagnostic information in amodal representations is to attach numeric weights such as (Bird (lays eggs 0.95) (flies 0.76)). However, this solution requires the semi-diagnostic features to be represented as fully as the diagnostic features. In many cases, it is desirable to have the quality or size of the actual feature representation correspond to its diagnosticity.

Detailed mechanisms are now available to describe the role of binding in perception. Elements from a person's left-eye image are bound to their corresponding elements from the right-eye image, such that globally harmonious structures, and the perception of depth, are established (Marr and Poggio, 1979). Similarly, elements from one 'frame' of an event are bound to their corresponding elements from the subsequent frame on the basis of local similarity and their respective roles within frames, in order to establish the perception of motion (Dawson, 1991). Finally, models of object recognition work by binding the parts of a perceived object to the parts of an internal object model (Hummel and Biederman, 1992). Thus, a suitcase is recognized by binding one part of an image to a 'handle' role and the rest of the image to a 'container' role. Bindings are created by first passing inhibitory and excitatory signals between image elements so that elements belonging to a single part are activated in phase with each other. Bindings, thus, can be implemented by synchronized patterns of activity between the elements. In the case of Hummel and Biederman's model, these bindings may connect amodal symbols such as the 'handle' role to a part, but also occur between perceptually-grounded representations such as two image parts.

While these models of perceptual binding offer mechanisms for the structuring of abstract thought (for an application to analogical reasoning, see Hummel and Holyoak, 1997), they also offer the exciting possibility of intermediate degrees of binding. This possibility is neglected by standard propositional representations because of their explicit, all-or-none assignment of arguments to roles. In contrast, most perceptually-motivated models of binding provide a dynamic and temporally extended mechanism for establishing structures, with completely unambiguous, one-to-one bindings as only a special case. For example, in Goldstone and Medin (1994; Goldstone, 1994a) model of similarity assessments, parts of two scenes are gradually placed into alignment to the extent that they are similar and play the same role in their respective scenes. Until the alignments are fully established, a particular part may be bound 70% to another part. Such an approach holds promise for abstract cognition as well. Even in the abstract domains of analogical reasoning and problem solving, the determination of abstract correspondences is imperfect (Gentner and Toupin, 1986), and is sensitive to superficially misleading features (Holyoak and Koh, 1987; Ross, 1987). In reading comprehension, fast processing is characterized by minimal inferences that can be established by unstructured priming rather than thoroughly worked out conceptual assignments (McKoon and Ratcliff, 1992). People's primitive preferences also seem to use imperfect bindings. Drinks labeled 'not poisonous' are rated as undesirable, presumably because 'poisonous' is somewhat free-floating, not completely bound to 'not' (Rozin et al., 1990). In sum, perception requires structured binding as much as does abstract cognition, and mechanisms of perceptual binding may better explain cases of intermediate degrees of binding than do amodal propositions.

3.6.4. Differentiation, subcategories and dimensions

A primary method for refining thought is to differentiate – to take a rough category and sub-divide it into smaller categories. Researchers frequently describe

the development of children's concepts as differentiation. Children, who originally confuse weight and density (Smith et al., 1985), or goats and dogs (Mandler and McDonough, 1991; Mandler et al, 1991), develop cognitively by creating distinct concepts for the originally confounded entities. Interestingly, child development also involves a parallel process of dimensionalization – breaking apart dimensions that were originally fused. Evidence suggests that dimensions easily separated by adults, such as the brightness and size of a square, are fused together for children (Kemler, 1983; Smith, 1989). For example, children have difficulty identifying whether two objects differ on their brightness or size, even though they can easily see that they differ in some way. Both differentiation and dimensionalization occur throughout one's lifetime. Tanaka and Taylor (1991) show that experts in a domain become increasingly adept at making subordinate-level categorizations. Dog experts, for example, can distinguish between German shepherd and golden retrievers as fast as they can distinguish between dogs and cats; non-experts are much faster at the latter discrimination. Likewise, dimensions that are psychologically fused for most adults, such as the chroma and hue of a color, can become separated with practice (Goldstone, 1994c). Artists and scientists who deal with colors regularly are better than non-experts at extracting dimensional information about chroma while ignoring hue differences (Burns and Shepp, 1988).

The primary reason to think that cognitive differentiation and perceptual dimensionalization share fundamental mechanisms is that it is simply hard to make a principled distinction between dimensions (features) and concepts. Many of the features that have been proposed for describing concepts (e.g. 'nests' and 'lays eggs' for 'bird,' 'has wheels' and 'has engine' for 'car' and 'sit on, and 'legs' for 'chair') are concepts in their own right (Schyns et al., 1998). If we dispense with the traditional division between concepts and dimensions, then we can potentially take advantage of the interesting work on computational mechanisms that ground conceptual differentiation in perceptual dimensionalization (de Sa and Ballard, 1997; Smith et al., 1997). For example, in the neural network model of Smith et al., the network's ability to assess specific similarities and differences between objects depends on the network's first creating separate perceptual dimensions along different units. Furthermore, this link also explains why the left hemisphere appears specialized for both fine perceptual discriminations, such as when stimuli have high spatial frequencies or when subjects are told to use a strict criteria for identification (Robertson and Lamb, 1991), and relatively differentiated language use, such as choosing the appropriate meaning of an ambiguous word (Burgess and Simpson, 1988). Conversely, the right hemisphere seems better tuned to relatively global, multi-dimensional perceptual aspects and relatively broad linguistic categories (Brownell et al., 1990). Thus, it is at least plausible to believe that the processes that split percepts into dimensions and more abstract categories into sub-categories may be related.

3.6.5. Cross-modal matching

In the spirit of speculative inquiry, we include two additional parallels between

conceptual and perceptual processes. An ability to reason analogically may borrow from perceptual processes that underlie cross-modal matching, which is characterized by people's naturally tendency to link distinct sensory domains. The most publicized example of this phenomenon is the synesthesia experienced by rare individuals who experience, for example, vivid auditory percepts when shown visual stimuli and vice versa. However, even normal adults seem predisposed to directly perceive relations between separate sensory modalities such as loudness, pitch and color (Melara, 1989).

3.6.6. Productivity

The operation of productivity (combining atomic units together to create larger structures) is important for abstract thought because it produces mundane creativity; a potentially infinite number of new thoughts can be generated by recombining existing thoughts in new arrangements. Although productivity is typically associated with amodal systems, perceptual representations can support an equivalent operation: spatio-temporal juxtaposition (Barsalou, 1993; Barsalou and Prinz, 1997; Prinz, 1997; Prinz and Barsalou, 1997). Separate images can be juxtaposed to produce new images. Furthermore, evidence exists that new interpretations can accompany these juxtapositions. When asked to rotate an imagined 'D' counterclockwise 90°, and then intersecting an imagined 'J' such that the top of the 'J' touches the top of the rotated 'D,' subjects often are able to reinterpret their concatenation as an umbrella (Finke et al., 1989). Fodor and Pylyshyn (1988) argue that amodal symbols are required for compositionality, which is required, in turn, for productivity; our point is that images, not merely amodal symbols, can and do function as the atomic units that are combined.

3.6.7. Summary

The significance of these parallels is two-fold. The first point is that many properties of abstract cognition, when explored from the perspective of processes that could furnish them, are also found in perceptual systems (for other parallels between perceptual and conceptual categories, see Medin and Barsalou, 1987). Claims that abstract cognition is special because it highlights relevant properties, is abstractive, has argument structure, permits analysis into components, or allows productivity are weakened by the presence of perceptual equivalents. In several cases, much more is known about the mechanisms of these perceptual processes than is known about their conceptual counterparts (Ullman, 1984), and, pragmatically speaking, we would be well-advised to use this knowledge to guide our understanding of abstract thought. The second, more speculative, point is that these parallels are hardly coincidental but arise because important mechanisms are common to perception and cognition. An extreme version of this position is unlikely, but an examination of individual differences, task manipulations and neuropsychological data provides enough evidence for correlations between perceptual and conceptual tasks to encourage exploration of the possibility that some of the process pairs of Table 1 are linked by process-overlap rather than by mere analogy.

4. Remaining issues

In an attempt to reunite perceptual and conceptual processing, we have argued that perceptual mechanisms provide unexpectedly rich and useful resources for implementing a conceptual system. First, perceptual representations are powerful, because they can implicitly represent properties of the external world in an analog fashion. Fairly abstract properties, such as the regression coefficient and minimal spanning distance, can be represented in analog systems without being calculated explicitly. Second, impressions of overall, undifferentiated similarity seem to be perceptually primary and to be exactly the perceptually-constrained type of similarity that is useful in creating many common categories. Third, perceptual similarity is not static, but changes as a function of the categorization demands confronting an organism. This flexibility reduces the gap between perception and sophisticated analytic concepts. Fourth, even when concepts seem to have little perceptual basis, their origins can often be traced to perceptual processing. Fifth, people appear to perform perceptual simulation in conceptual tasks that have no perceptual demands. Sixth, striking commonalities exist between the mechanisms that process abstract information and those established for perception. Correlations between cognitive tasks (e.g. problem solving, language and reasoning) and perceptual tasks suggest that shared mechanisms underlie them.

4.1. The perceptual/conceptual distinction

Is there a continuum from perceptual to conceptual representations, and, if so, what varies along this continuum? Given the top-down influences of concepts on perception (e.g. Goldstone, 1995b) and the surprisingly far-reaching influences of expectations on low-level perceptual judgments (e.g. Peterson and Gibson, 1994), searching for *the* boundary between perception and conception is most likely futile. However, it may be useful to describe a *continuum* from perceptual to conceptual. What varies along this continuum is how much and what sort of top-down processing has been done to bottom-up input information. Specifying exactly where expectations and conceptual pressures influence processing along the perception/conception continuum is a real, although highly empirical, question. The general principle that conceptual processes can more readily be tuned to particular demands than perceptual processes probably has some validity. When we categorize 'T's and tilted 'T's as belonging to the same letter category, it increases their rated similarity (here, similarity ratings, despite 'similarity' in their name, are relatively conceptual tasks!). However, categorizing 'T's and tilted 'T's has little influence on more perceptual measures of similarity, such as our ability to quickly spot the borderline between a group of 'T's and a group of tilted 'T's (Beck, 1966). Similarly, one's momentary goals have clear and large influences on cognitive processes such as description and inference (Barsalou and Sewell, 1984), but less influence on perceptual processes such as figure/ground segmentation, color afterimages, edge detection, and same/different discrimination, (or it is more surprising when they do have an influence (Moscovici and Personnaz, 1991)).

We also assume that conception differs from perception in degree of productivity. In conception, one can combine perceptual representations in ways that go far beyond perception. In imagining the Cheshire Cat from *Alice in Wonderland*, productivity is freed from the constraints of actual perception. Real cats don't have human smiles, and their bodies don't fade in and out while their smiles remain. Although conception may have fundamental underpinnings in perception, its ability to manipulate schematic perceptual representations productively allows it to go considerably further.

4.2. The perceptual/abstract distinction

We are less sanguine about the usefulness of a perceptual/abstract continuum. Perception often involves the abstraction of certain elements. Concrete details are often blurred over, ignored or actively suppressed. Similarly, selective attention can focus analytically on a perceptual dimension, extracting information about that aspect of experience. As Arnheim (1969) points out, the difference between realistic and abstract art is not one of concreteness or relevance of perception. Abstract art is, of course, concrete. An understanding of perceptual processes is often of fundamental importance for creating and appreciating abstract art. The abstractions at work in a piece of music or painting are often not the same sort of abstractions present in a novel. For example, in a painting, abstractions often deal with spatial relations, relations between colors and the manner of dividing and integrating different areas on the canvas. Often, it is notoriously difficult to verbally describe these perceptual abstractions (e.g. those in the works of Picasso). The representations in a particular perceptual domain (hearing or seeing) may be highly abstract even though they are constrained by concrete qualities of the particular domain.

A corollary to the notion that abstractions can be abstract even when tied to their particular perceptual domain, is that it is misleading to equate 'perceptual' and 'superficial.' A traditional assumption, particularly in research on analogical reasoning, is that comparisons based on perceptual aspects are superficial. Accordingly, the *deep* similarity between time bombs and cigarettes is that they both cause delayed damage; the fact that they both involve fire is deemed superficial (and hence less interesting or important). In contrast, following Bassok (Bassok and Olseth, 1995; Bassok et al., 1995), we believe that perceptual aspects typically are at least cues for the abstraction that is built, and often are never removed from the abstraction. The similarity between diving boards and bed springs depends on perceptual aspects (a bouncing motion with gradual deceleration and acceleration). These perceptual aspects are 'deep' in the sense that they permit widespread causal inferences between highly dissimilar objects, and that they result from general physical laws. Many physical, biological and psychological principles that are discussed in everyday life produce perceptual effects, and our perceptual systems have been refined to make this apparent. Given this, perceptually-based comparisons probably yield more reliable inferences than those produced by analogies that are completely stripped of their perceptual grounds.

4.3. Concluding remarks

In endorsing the perceptual/conceptual distinction, we essentially argue that being conceptual is being abstract. In not endorsing the perceptual/abstract distinction, we implicitly deny the requirement that abstract knowledge be amodal. Instead, we argue that abstract knowledge can be constituted from perceptual bases.

There is a certain tension between two of our arguments. On the one hand, we have argued that perceptual representations are useful because of the constraints and mechanisms inherent in their unprocessed, image-like representations. On the other hand, we have also argued that perceptual systems provide abstractive and selective processes that allow perception to distort or transform sensory inputs. These positions can be reconciled. Perceptual processes vary in how much they transform the original sensory input, yet concrete aspects of the input are rarely, if ever, discarded completely. The special case of completely amodal symbols can be approximated by highly-transformed perceptual representations. There is a continuum of distance from sensory input, with amodal symbols constituting an ideal endpoint. Recognizing this continuum allows us to imagine thought processes that highlight, structure and abstract information, without discarding all perceptual contents. Rather than implicating dichotomies between bound versus unbound representations, holistic versus analytic processes and abstract versus concrete thought, the reviewed evidence suggests continua instead. In developing future theories of cognition, it may be fruitful to design architectures capable of implementing a wide variety of special cases along these continua, as required under varying task demands.

We are left with the question of what role, if any, do amodal symbols play in conceptual knowledge? We have seen that perceptual mechanisms can accomplish many of the functions that are well-known for amodal symbol systems. On the basis of such observations, one might be inclined to adopt the eliminativist view that amodal symbols are unnecessary. We can think of two reasons why one might want to maintain amodal symbols. First, if there is some necessary function of intelligence that amodal symbols can accomplish that perceptual ones cannot, then this is an obvious reason for maintaining amodal symbols. Second, we might find that amodal symbols deliver certain conceptual functions more efficiently that perceptual symbols. Even though perceptual symbols could implement the same functions, amodal symbols may do so more optimally.

Regardless of where one comes down on the eliminative/agnostic distinction, we believe that perceptual mechanisms underlie conceptual processing to a considerable degree. From overall similarity to analytic rules, many sources of evidence implicate perception in conception.

Acknowledgements

References

Allen, S.W., Brooks, L.R., 1991. Specializing the operation of an explicit rule. Journal of Experimental Psychology: General 120, 3–19.

Anderson, J.A., 1997. Seven times seven is about 50: learning arithmetic with a neural network. In: S. Sternberg (Ed.), Invitation to Cognitive Science. MIT Press, Cambridge, MA, in press.

Anderson, M.C., Spellman, B.A., 1995. On the status of inhibitory mechanisms in cognition: memory retrieval as a model case. Psychological Review 102, 68–100.

Arnheim, R., 1969. Visual Thinking. Faber and Faber, London.

Ashby, F.G., 1992. Multidimensional Models of Perception and Cognition. Erlbaum, Hillsdale, NJ.

Barsalou, L.W., 1991. Deriving categories to achieve goals. In: Bower, G.H. (Ed.), The Psychology of Learning and Motivation: Advances in Research and Theory, Vol. 27. Academic Press, New York.

Barsalou, L.W., 1992. Frames, concepts, and conceptual fields. In: Kittay, E., Lehrer, A. (Eds.), Frames, Fields, and Contrasts: New Essays in Semantic and Lexical Organization. Erlbaum, Hillsdale, NJ, pp. 21–74.

Barsalou, L.W., 1993. Flexibility, structure, and linguistic vagary in concepts: manifestations of a compositional system of perceptual symbols. In: Collins, A.C., Gathercole, S.E., Conway, M.A., (Eds.), Theories of Memories. Erlbaum, London, pp. 29–101.

Barsalou, L.W., Hale, C.R., 1993. Components of conceptual representation: from feature lists to recursive frames. In: Van Mechelen, I, Hampton, J., Michalski, R, Theuns, P (Eds.), Categories and Concepts: Theoretical Views and Inductive Data Analysis. Academic Press, San Diego, CA, pp. 97–144.

Barsalou, L.W., Prinz, J.J., 1997. Mundane creativity in perceptual symbol systems. In: Ward, T.B., Smith, S.M., Vaid, J. (Eds.), Creative Thought: An investigation of Conceptual Structures and Processes. American Psychological Association, Washington, DC, pp. 267–307.

Barsalou, L.W., Sewell, D.R., 1984. Constructing categories from different points of view. Emory Cognition Report No. 2. Emory University, Atlanta, GA.

Barsalou, L.W., Solomon, K.O., Wu, L.L., 1997. Perceptual simulation in conceptual tasks. In: Hiraga, M.K., Sinha, C, Wilcox, S. (Eds.), Cultural, Typological, and Psychological Perspectives in Cognitive Linguistics: the Proceedings of the 4th Conference of the International Cognitive Linguistics Association, Vol. 3. John Benjamins, Amsterdam, pp. 337–346.

Barsalou, L.W., Yeh, W., Luka, B.J., Olseth, K.L., Mix, K.S., Wu, L., 1993. Concepts and meaning. In: Beals, K, Cooke, G., Kathman, D., McCullough, K.E., Kita, Testen, D. (Eds.), Chicago Linguistics Society 29: Papers from the Parasession on Conceptual Representations. Chicago Linguistics Society, University of Chicago, pp. 23–61.

Barwise, J., Etchemendy, J., 1990. Information, infons, and inference. In: Cooper, R., Mukai, K., Perry, J. (Eds.), Situation Theory and its Applications. University of Chicago Press, Chicago, pp. 33–78.

Barwise, J., Etchemendy, J., 1991. Visual information and valid reasoning. In: Zimmerman, W., Cunningham, S. (Eds.), Visualization in Mathematics. Mathematical Association of America, Washington, pp. 9–24.

Bassok, M., Olseth, K., 1995. Object-based representations: transfer between cases of continuous and discrete models of change. Journal of Experimental Psychology: Learning, Memory, and Cognition 21, 1522–1538.

Bassok, M., Wu, L., Olseth, K., 1995. Judging a book by its cover: interpretative effects of content on problem-solving transfer. Memory and Cognition 23, 354–367.

Bearse, M.A., Freeman, R.D., 1994. Binocular summation in orientation discrimination depends on stimulus contrast and duration. Vision Research 34, 19–29.

Beck, J., 1966. Effect of orientation and of shape similarity on perceptual grouping. Perception and Psychophysics 1, 300–302.

Beech, A., Powell, T., McWilliam, J., Claridge, G., 1989. Evidence of reduced 'cognitive inhibition' in schizophrenia. British Journal of Clinical Psychology 28, 109–116.

Biederman, I., Shiffrar, M.M., 1987. Sexing day-old chicks: a case study and expert systems analysis of a difficult perceptual-learning task. Journal of Experimental Psychology: Learning, Memory, and Cognition 13, 640–645.

Brooks, L.R., 1978. Non-analytic concept formation and memory for instances. In: Rosch, E., Lloyd, B.B. (Eds.), Cognition and Categorization. Erlbaum, Hillsdale, NJ. pp. 169–211.

Brownell, H.H., Simpson, T.L., Bihrle, A.M., 1990. Appreciation of metaphoric alternative word meanings by left and right brain-damaged patients. Neuropsychologica 28, 375–383.

Burgess, C., Simpson, G.B., 1988. Cerebral hemispheric mechanisms in the retrieval of ambiguous word meanings. Brain and Language 33, 86–103.

Burns, B., 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., 1985. Conceptual Change in Childhood. MIT Press, Cambridge, MA.

Chi, M.T.H., Feltovich, P., Glaser, R., 1981. Categorization and representation of physics problems by experts and novices. Cognitive Science 5, 121–152.

Dawson, M.R., 1991. The how and why of what went where in apparent motion: modeling solutions to the motion correspondence problem. Psychological Review 98, 569–603.

Dewdney, A.K., 1985. Analog gadgets that solve a diversity of problems and raise an array of questions. Scientific American, June, 18–29.

de Sa, V., Ballard, J., 1997. Learning Perceptual Dimensions. In: Medin, D.L., Goldstone, R.L., Schyns, P. (Eds.) Psychology of Learning and Motivation. Academic Press, San Diego, pp. 309–352.

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

Farah, M.J., 1990. *V*isual Agnosia: Disorders of Object Recognition and What they tell us about Normal Vision. MIT Press, Cambridge, MA.

Finke, R.A., 1986. Mental imagery and the visual system. Scientific American, March, 98–104.

Finke, R.A., Pinker, S., Farah, M., 1989. Reinterpreting visual patterns in mental imagery. Cognitive Science 13, 51–78.

Foard, C.F., Kemler Nelson, D.G., 1984. Holistic and analytic modes of processing: the multiple determinants of perceptual analysis. Journal of Experimental Psychology: General 113, 94–111.

Fodor, J., 1975. The Language of Thought. Crowell, New York.

Fodor, J.A., Pylyshyn, Z.W., 1988. Connectionism and cognitive architecture: a critical analysis. Cognition 28, 3–71.

Garner, W.R., 1974. The Processing of Information and Structure. Wiley, New York.

Garner, W.R., 1978. Aspects of a stimulus: features, dimensions, and configurations. In: Rosch, E., Lloyd, B.B. (Eds.), Cognition and Categorization. Erlbaum, Hillsdale, NJ.

Gelman, S.A., 1988. The development of induction within natural kind and artifact categories. Cognitive Psychology 20, 65–95.

Gelman, S.A., Markman, E.M., 1986. Categories and induction in young children. Cognition 23, 183–209.

Gelman, R., 1990. First principles organize attention to and learning about relevant data: number and the animate-inanimate distinction as examples. Cognitive Science 14, 79–106.

Gentner, D., Toupin, C., 1986. Systematicity and surface similarity in the development of analogy. Cognitive Science 10, 277–300.

Gernsbacher, M.A., Faust, M.E., 1991. The mechanism of suppression: a component of general comprehension. Journal of Experimental Psychology: Learning, Memory, and Cognition 17, 245–262.

Gick, M.L., Holyoak, K.J., 1983. Schema induction and analogical transfer. Cognitive Psychology 15, 1–39.

Lakoff, G., 1986. Women, Fire and Dangerous Things: what Categories tell us about the Nature of Thought. University of Chicago Press, Chicago.

Glucksberg, S., Keysar, B., 1990. Understanding metaphorical comparisons: beyond similarity. Psychological Review 97, 3–18.

Goldstone, R.L., 1994a. Similarity, interactive activation, and mapping. Journal of Experimental Psychology: Learning, Memory, and Cognition 20, 3–28.

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

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

Goldstone, R.L., 1995a. Mainstream and avant-garde similarity. Psychologica Belgica 35, 145–165.

Goldstone, R.L., 1995b. Effects of categorization on color perception. Psychological Science 6, 298–304.

Goldstone, R.L., Medin, D.L., 1994. The time course of comparison. Journal of Experimental Psychology: Learning, Memory, and Cognition 20, 29–50.

Goldstone, R.L., Schyns, P., 1994. Learning new features of representation. Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society. Erlbaum, Hillsdale, New Jersey, pp. 974–978.

Gould, S.J., 1993. Eight Little Piggies. Norton, New York.

Goodman, N., 1972. Seven strictures on similarity. In: Goodman, N. (Ed.), Problems and Projects. Bobbs-Merrill, New York, pp. 35–41.

Hadamard, J., 1949. The Psychology of Invention in the Mathematical Field. Dover Books, New York.

Harnad, S., 1987. Categorical Perception. Cambridge University Press, Cambridge.

Holyoak, K.J., Koh, K., 1987. Surface and structural similarity in analogical transfer. Memory and Cognition 15, 332–340.

Hummel, J.E., Biederman, I., 1992. Dynamic binding in a neural network for shape recognition. Psychological Review 99, 480–517.

Hummel, J.E., Holyoak, K.J., 1997. Distributed representations of structure: a theory of analogical access and mapping. Psychological Review 104, 427–466.

Huttenlocher, J., Jordan, N., Levine, S., 1994. A mental model for early arithmetic. Journal of Experimental Psychology: General 123, 284–296.

Johnston, W.A., Dark, V.J., 1986. Selective attention. Annual Review of Psychology 37, 43–75.

Johnson-Laird, P.N., 1983. Mental Models. Harvard University Press, Cambridge, MA.

Jones, S.S., Smith, L.B., 1993. The place of perception in children's concepts. Cognitive Development 8, 113–140.

Kant, I., 1965. The Critique of Pure Reason (N.K. Smith, Trans.). St. Martin's Press, New York. (Original work published in 1787.).

Karni, A., Sagi, D., 1991. Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. Proceedings of the National Academy of Sciences of the United States of America 88, 4966–4970.

Keil, F.C., 1989. Concepts, Kinds and Development. Bradford Books/MIT Press, Cambridge, MA.

Kelly, M.H., Keil, F.C., 1987. Metaphor comprehension and knowledge of semantic domains. Metaphor and Symbolic Activity 2, 33–51.

Kemler, D.G., 1983. Holistic and analytic modes in perceptual and cognitive development. In: Tighe, T.J., Shepp, B.E. (Eds.), Perception, Cognition, and Development: Interactional Analyses. Erlbaum, Hillsdale, NJ, pp. 77–101.

Kirk, G.S., Raven, J.E., 1962. The Presocratic Philosophers. Cambridge University Press, Cambridge.

Kosslyn, S.M., 1976. Can imagery be distinguished from other forms of conceptual representation? Evidence from studies of information retrieval times. Memory and Cognition 4, 291–297.

Kosslyn, S.M., 1980. Image and Mind. Harvard University Press, Cambridge, MA.

Kroska, A., Goldstone, R.L., 1996. Dissociations in the similarity and categorisation of emotions. Cognition and Emotion 10, 27–45.

Langacker, R.W., 1986. An introduction to cognitive grammar. Cognitive Science 10, 1–40.

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

Lenat, D.B., Guha, R.V., 1994. Enabling agents to work together. Communications of the ACM 37, 203–215.

Liotti, M., Dazzi, S., Umilta, C., 1993. Deficits of the automatic orienting of attention in schizophrenic patients. Journal of Psychiatric Research 27, 119–130.

Locke, J., 1959. An essay concerning human understanding, first ed. (Vols. I and II). Dover, New York. (Original work published in 1690.).

Logan, G.D., Etherton, J.L., 1994. What is learned during automatization? The role of attention in constructing an instance. Journal of Experimental Psychology: Learning, Memory, and Cognition 20, 1022–1050.

Lovass, O.I., Koegel, R.L., Schreibman, L., 1979. Stimulus overselectivity in autism: a review of research. Psychological Bulletin 86, 1236–1254.

Mandler, J.M., 1992. How to build a baby: II. Conceptual primitives. Psychological Review 99, 587–604.

Mandler, J.M., Bauer, P.J., McDonough, L., 1991. Separating the sheep from the goats: differentiating global categories. Cognitive Psychology 23, 263–298.

Mandler, J.M., McDonough, L., 1991. Concept formation in infancy. Cognitive Development 8, 291–318.

Markman, A.B., Gentner, D., 1993. Structural alignment during similarity comparisons. Cognitive Psychology 25, 431–467.

Marr, D., Poggio, T., 1979. A computational theory of human stereo vision. Proceedings of the Royal Society of London 204, 301–328.

McKoon, G., Ratcliff, R., 1992. Inference during reading. Psychological Review 99, 440–466.

Medin, D.L., 1989. Concepts and conceptual structure. American Psychologist 44, 1469–1481.

Medin, D.L., Barsalou, L.W., 1987. Categorization processes and categorical perception. In: Harnad, S., (Ed.) Categorical Perception. Cambridge University Press, Cambridge, pp. 455–490.

Medin, D.L., Goldstone, R.L., Gentner, D., 1993. Respects for similarity. Psychological Review 100, 254–278.

Medin, D.L., Ortony, A., 1989. Psychological essentialism. In: Vosniadou, S., Ortony, A. (Eds.), Similarity and Analogical Reasoning. Cambridge University Press, Cambridge, pp. 179–195.

Melara, R.D., 1989. Similarity relations among synesthetic stimuli and their attributes. Journal of Experimental Psychology: Human Perception and Performance 115, 212–231.

Melara, R.D., Marks, L.E., 1990. Dimensional interactions in language processing: investigating directions and levels of crosstalk. Journal of Experimental Psychology: Learning, Memory, and Cognition 16, 539–554.

Miller, A., 1987. Cognitive styles: an integrated Model. Educational Psychology 7, 251–268.

Moscovici, S., 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.

Mullin, P.A., Egeth, H.E., 1989. Capacity limitations in visual word processing. Journal of Experimental Psychology: Human Perception and Performance 15, 111–123.

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

Murphy, G.L., Spalding, T., 1995. Knowledge, similarity, and concept formation. Psychologica Belgica 35, 127–144.

Neill, W.T., 1977. Inhibition and facilitation processes in selective attention. Journal of Experimental Psychology: Human Perception and Performance 3, 444–450.

Neisser, U., Becklen, R., 1975. Selective looking: attending to visually specified events. Cognitive Psychology 7, 480–494.

Nickerson, R.S., 1972. Binary Classification reaction time: a review of some studies of human information-processing capabilities. Psychonomic Monograph Supplements 4 (6), 275–317.

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

Nosofsky, R.M., 1992. Exemplar-based approach to relating categorization, identification, and recognition. In: Ashby, F.G. (Ed.), Multidimensional Models of Perception and Cognition. Erlbaum, Hillsdale, NJ, pp. 207–225.

Olseth, K.L., Barsalou, L.W., 1995. The spontaneous use of perceptual representations during conceptual processing. Proceedings of the seventeenth annual meeting of the Cognitive Science Society. Erlbaum, Hillsdale, NJ, pp. 310–315.

Ortony, A., 1979. Beyond literal similarity. Psychological Review 86, 161–180.

Palmer, S.E., 1978. Structural aspects of visual similarity. Memory and Cognition 6, 91–97.

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

Prinz, J.J., 1997. Perceptual Cognition: an Essay on the Semantics of Thought. Dissertation, Department of Philosophy, University of Chicago.

Prinz, J.J., Barsalou, L.W., 1997. Acquisition and productivity in perceptual symbol systems: an account of mundane creativity. In: Dartnall, T.H. (Ed.), Creativity, Computation, and Cognition. MIT/AAAI Press, Cambridge, MA, pp. 231–251.

Quine, W.V.O., 1977. Natural kinds. In: Schwartz, S.P. (Ed.), Naming, Necessity, and Natural Kinds. Cornell University Press, Ithaca, NY, pp. 35–59.

Rips, L.J., 1986. Mental muddles. In: Brand, M., Harnish, R.M. The Representation of Knowledge and Belief. University of Arizona Press, Tuscon, pp. 258–286.

Rips, L.J., 1989. Similarity, typicality, and categorization. In: Vosniadu, S., Ortony, A. (Eds.), Similarity, Analogy, and Thought. Cambridge University Press, Cambridge, pp. 21–59.

Rips, L.J., Collins, A., 1993. Categories and resemblances. Journal of Experimental Psychology: General 122, 468–486.

Robertson, L.C., Lamb, M.R., 1991. Neuropsychological contributions to theories of part/whole organization. Cognitive Psychology 23, 299–330.

Rock, I., Gutman, D., 1981. The effect of inattention on form perception. Journal of Experimental Psychology: Human Perception and Performance 7, 275–185.

Rosch, E., 1975. Cognitive representations of semantic categories. Journal of Experimental Psychology: Human Perception and Performance 1, 303–322.

Rosch, E., Mervis, C.B., 1975. Family resemblance: studies in the internal structure of categories. Cognitive Psychology 7, 573–605.

Rosch, E., Mervis, C.B., Gray, W., Johnson, D., Boyes-Braem, P., 1976. Basic objects in natural categories. Cognitive Psychology 7, 573–605.

Ross, B.H., 1987. This is like that: the use of earlier problems and the separation of similarity effects. Journal of Experimental Psychology: Learning, Memory, and Cognition 13, 629–639.

Ross, B.H., Perkins, S.J., Tenpenny, P.L., 1990. Reminding-based category learning. Cognitive Psychology 22, 460–492.

Rozin, P., Markwith, M., Ross, B., 1990. The sympathetic magical law of similarity, nominal realism and neglect of negatives in response to negative labels. Psychological Science 1, 383–384.

Russell, B., 1956. Logic and Knowledge. Routledge, London. (Original work published in 1919.).

Schreiber, H., Stolz-Born, G., Kornhuber, H.H., Born, J., 1992. Even before cognitive effects emerge, there are visual differences: event-related potential correlates of impaired selective attention in children at high risk for schizophrenia. Biological Psychiatry 32, 634–651.

Schwartz, S., 1982. Is there a schizophrenic language? Behavioral and Brain Sciences 5, 579–626.

Schyns, P.G., Goldstone, R.L., Thibaut, J-P., 1998. The development of features in object concepts. Behavioral and Brain Sciences (in press).

Searle, J.R., 1980. Minds, brains, and programs. Behavioral and Brain Sciences 3, 417–424.

Sekuler, R.W., Abrams, M., 1968. Visual sameness: a choice time analysis of pattern recognition. Journal of Experimental Psychology 77, 232–238.

Shanks, D.R., St. John, M.F., 1994. Characteristics of dissociable human learning systems. Behavioral and Brain Sciences 17, 367–447.

Shanon, B., 1988. On similarity of features. New Ideas in Psychology 6, 307–321.

Shepard, R.N., 1984. Ecological constraints on internal representations: resonant kinematics of perceiving, imagining, thinking, and dreaming. Psychological Review 91, 417–447.

Shaw, M., 1982. Attending to multiple sources of information: I. The integration of information in decision making. Cognitive Psychology 14, 353–409.

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

Shiffrin, R.M., Schneider, W., 1977. Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. Psychological Review 84, 127–190.

Sloman, S.A., 1996. The empirical case for two systems of reasoning. Psychological Bulletin 119, 3–22.

Smith, C., Carey, S., Wiser, M., 1985. On differentiation: a case study of the development of the concepts of size, weight and density. Cognition 21, 177–237.

Smith, E.E., Langston, C., Nisbett, R., 1992. The case for rules in reasoning. Cognitive Science 16, 1–40.

Smith, E.E., Sloman, S.A., 1994. Similarity- versus rule-based categorization. Memory and Cognition 22, 377–386.

Smith, L.B., 1989. From global similarity to kinds of similarity: the construction of dimensions in development. In: Vosniadou, S., Ortony, A. (Eds.), Similarity and Analogical Reasoning. Cambridge University Press, Cambridge, pp. 146–178.

Smith, L.B., Gasser, M., Sandhofer, C. 1997. Learning to talk about the properties of objects: a network model of the development of dimensions. In: Goldstone, R.L., Schyns, P.G., Medin, D.L. (Eds.) Psychology of Learning and Motivation, Vol. 36. Academic Press, San Diego, CA, pp. 219–256.

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

Smith, J.D., Kemler Nelson, D.G., 1984. Overall similarity in adults' classification: the child in all of us. Journal of Experimental Psychology: General 113, 137–159.

Solomon, K.O., 1997 The Spontaneous use of Perceptual Representations during Conceptual Processing. Dissertation, Department of Psychology, University of Chicago.

Spelke, E.S., 1990. Principles of object perception. Cognitive Science 14, 29–56.

Stigler, J.W., 1984. 'Mental abacus': the effect of abacus training on Chinese children's mental calculation. Cognitive Psychology 16, 145–176.

Talmy, L., 1983. How language structures space. In: Pick, H., Acredelo, L (Eds.), Spatial Orientation: Theory, Research, and Application. Plenum Press, New York, pp. 225–282.

Tanaka, J.W., Taylor, M., 1991. Object categories and expertise: is the basic level in the eye of the beholder? Cognitive Psychology 23, 457–482.

Tipper, S.P., 1992. Selection for action: the role of inhibitory mechanisms. Current Directions in Psychological Science 1, 105–109.

Triesman, A.M., 1960. Contextual cues in selective attention. Quarterly Journal of Experimental Psychology 12, 242–248.

Treisman, A.M., Gelade, G., 1980. A feature-integration theory of attention. Cognitive Psychology 12, 97–136.

Tversky, A., 1977. Features of similarity. Psychological Review 84, 327–352.

Ullman, S., 1984. Visual routines. Cognition 18, 97–159.

Wapner, S., Demick, J., 1991. Field Dependence-Independence: Cognitive Style across the Life Span. Erlbaum, Hillsdale, NJ.

Ward, T.B, 1983. Response tempo and separable-integral responding: evidence for an integral-to-separable processing sequence in visual perception. Journal of Experimental Psychology: Human Perception and Performance 9, 103–112.

Wu, L., 1995. Perceptual Representation in Conceptual Combination. Doctoral dissertation, University of Chicago.