The Label Feedback Hypothesis:
Linguistic Influences on Visual Processing.

by

Gary Lupyan

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

**Doctor of Philosophy
in
Psychology**

**Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213.**

August, 2007

Abstract

Humans are the only animals to have names for their categories. This makes linguistic communication possible. Additionally, category names (i.e., words) have been argued to stabilize abstract ideas in working memory (Clark, 1998) and make them available for inspection (Clark & Karmiloff-Smith, 1993; Rumelhart, Smolensky, McClelland, & Hinton, 1986; Vygotsky, 1962). Few studies, however, have tested such ideas directly, and the dominant view is that language, rather than shaping and modulating mental representations, is simply a tool for reporting them (Bloom & Keil, 2001; Gleitman & Papafragou, 2005; Li & Gleitman, 2002; Pinker, 1995).

This thesis provides support for the hypothesis that rather than simply being the outputs of "alinguistic" conceptual representations, verbal labels participate in creating and altering them. I review evidence from language development and language deficits for the tight linkage between language and categorization, and summarize previous work showing that learning entirely redundant category names facilitates the acquisition of novel categories, that labeling familiar categories leads to poorer memory for specific exemplars, and that verbal interference produces categorization impairments of the kind found in aphasic patients.

This thesis extends these findings by investigating on-line effects of labels on visual perception using the paradigm of visual search. Hearing familiar labels is found to facilitate visual processing as revealed by faster visual search performance. This facilitation is observed when the target or the non-targets are labeled, and is mediated by the typicality of the labeled items. The facilitation following the labels is observed in relation to trials in which participants know the identity of all the visual elements, but do not actually hear the label. Thus, the observed effect of hearing labels is not due directly to the information they provide.

The experimental results are embedded within a new framework—The Label Feedback Hypothesis (LFH). According to the LFH, in learning to associate category names with the entities to which they refer, the label becomes associated with features most typical and distinctive to the named category. When activated, the activity produced by the label feeds down to modulate lower-level perceptual representations augmenting the bottom-up activity evoked by the visual stimulus with top-down information from the broader category. Because object names denote categories, naming an object induces its representation to become more categorical and less idiosyncratic, possibly resulting in the perceptual grouping of stimuli that match the named category. The framework developed in this thesis promises to be highly useful for understanding the often contradictory findings from studies aiming to understand effects of language on cognitive processing and the consequences of language deficits.

List of Tables

List of Figures

discriminating a B from a non-B. Bars show within-subject 95% confidence intervals. Asterisks signify significant differences between means at $p<.05$.

Acknowledgments

I am very grateful to my supervisor Jay McClelland for guidance, impeccable experimental intuitions, and for giving me the freedom to pursue this work. My committee—David Rakison and Anna Fisher—have come to my rescue on numerous occasions and I thank them for their support.

The core question that has driven much of my research—what do words do?—was first formulated while I was a research assistant in Morten Christiansen's lab at Cornell University. It was Morten's guidance that first led me down this path of science (and it has been a good one thus far).

I also thank Marlene Behrmann for turning me on to vision science, and David Plaut for his insights into connectionist modeling. My ideas about the importance of recurrence owe their origin in part to conversing with and reading the work of Mike Spivey.

This work has benefited enormously from discussions with my fellow graduate students past and present, including Dwight Kravitz, Cibu Thomas, Miriam Rosenberg-Lee, Evan Smith, and Erin Ingvalson.

For financial support, I thank the NSF Graduate Student Fellowship program, the Center for the Neural Basis of Cognition, and Jay McClelland's generosity. The travel made possible by this financial support has made a significant impact on this research by exposing me to the larger scientific community.

The data collection undertaken for this thesis (as well as experiments presented in other publications) would not be possible without my research assistants over the past few years: Brian Mathias, Ashleigh Molz, Bonnie Lee, Josh Lord, Rebecca Snyder, Leah Russel, Cheryl Gering, Lauren Winchester, and Karen Doersch.

Finally, I thank my entire family for their love and support, and my fiancée Ginger McCall for her love, help in times of need, and undying patience.

1 Aims and Motivations of the Thesis

What does language do? According to traditional accounts, language is "merely the formal and expressive medium that is [used] to describe mental representations" (Li & Gleitman, 2002 p . 290), that is, it is an output system—a way to communicate our ideas, which can be converted to language, but are themselves independent of it. Contrasting sharply with this conception of language is a view that language is "potentially catalytic and transformative of cognition" (Bowerman & Choi, 2001 p. 13).

This view of language as a cognitive "tool" (Clark, 1998) contrasts sharply with the traditional view which is also associated with the "cognitive priority hypothesis" (see section 2.3). Viewing language as an output system certainly does not preclude words from changing mental content. If they did not, linguistic communication would be impossible. The reason it is difficult to avoid thinking of a pink elephant when one tells you not to is that the word "elephant" denotes elephants. Linguistic comprehension requires the brain to construct a representation of the linguistic message. This simple fact is monumentally important, for it means that humans can, with simple utterances, activate clusters of associated features in the brains of their interlocutors. By saying "elephant" I have produced a representation of an elephant in your brain in the total absence of your encountering an actual elephant. Given a scattering of some objects, I can instruct you to find the red one, or the square one, or the red-square. I can also instruct you to ignore the highly salient red square and instead focus on the much less-salient gray one. Though performing such feats is not limited to humans; the ability to flexibly provide and comprehend such instructions is. I will argue that the use of language in such cases is not just a matter of communicative convenience. Rather,

> Language reflects and facilitates our most pervasive, open-ended manifestations of cognitive flexibility. The basic function of language *is* fast, flexible production and reconstruction of a practically unlimited range of selectively sculpted mental representations. No other behavior system in nature matches this potential for flexible representation (Deak, 2003 p. 318).

The goal of this thesis is to demonstrate that by virtue of the learned associations between words and their referents, words participate in the creation of categories they denote, and function on-line to selectively shape the perceptual representations that underlie our conceptual knowledge. Rather than simply a system for communicating to others one's pre-existing conceptual representations, I shall argue that verbal labels participate in constructing and modulating these representations. In addition to changing representations in our interlocutors, language changes representations in our own brains.

The theoretical component of this thesis develops a neurally inspired framework—the *Label Feedback Hypothesis* (LFH)—for understanding the disparate and often-conflicting findings of interactions between language, categorization, and visual processing. Following a description of the LFH I present a number of experiments that test the predictions of the framework in the domain of visual processing.

1.1 Critical Definitions

1.1.1 What are categories and what are they for?

In a complex environment an organism never encounters an exact stimulus more than once, and so perceiving a stimulus as an instance of a larger category—food, predator, mate—allows the organism to enact an appropriate response. This act of seeing something as X rather than simply seeing it (Wittgenstein, 1953), is fundamentally an act of categorization (Goldstone, Steyvers, Spencer-Smith, & Kersten, 2004), making categorization one of the keystones of cognition (Harnad, 2005).

One familiar definition of categorization is the process by which discriminable stimuli come to have common responses. This behaviorist definition has been modified in contemporary cognitive psychology to invoke mental representations: we say that members of a category share a representation and due to this common representation, these physically diverse stimuli produce both a common response and a shared concept. Notice that because stimuli from a common category share certain properties, placing a novel stimulus into a familiar category enables an organism to infer unobserved properties. For example, recognizing the head of an animal as being a dog's head allows an observer to infer what the rest of the animal looks like, how the animal will move, what the animal will sound like (e.g., Gelman & Coley, 1991) as well as what the appropriate response should be—is this the kind of thing I should pet or run away from?

Since even the same object encountered in different contexts, at different times, by an ever-changing nervous system, will produce different physical states, having concepts requires extracting some amount of invariance from the perceptual stream. Representing this invariance enables the detection of recurring events and requires that the representation of the event highlight some properties while abstracting over others (Harnad, 2005). Importantly, the process of abstraction need not, and (at least in humans) cannot be, fixed. The need to categorize an object flexibly depending on context and on the desired level of specificity (human / male / young / Bob / my brother) requires that different features of a representation be made salient. For humans in particular, the ability to form ad-hoc categories such as "things to take out of the house in case of a fire" clearly requires dynamically reconceptualizing which features of an object are relevant in a given task (Barsalou, 1983). This conception of categorization conflicts with the notion of categorization as a bottom-up "sieve" (at least if one assumes that the sieve is "fixed"). As I will argue in more detail, categorization is more usefully conceptualized as a flexible interplay between top-down and bottom-up processes. In the framework of the Label Feedback Hypothesis, verbal category labels induce categorical representations through top-down feedback (Section 3).

1.1.2 What are Labels?

While categorization is ubiquitous among animals, only humans have *names* for their categories. I refer to these as "labels." A label is simply the verbal response associated with a stimulus class (a label produced through "inner speech" is still a label despite not being verbalized overtly). Although the same stimulus can be labeled at multiple levels: dog, animal, furry, etc., here, I focus here on basic-level names—defined as the level at which classification is most "natural" (Rosch, 1973). Thus, the label for the dog concept is "dog" and the label for "2" is "two" (rather than "number" or "even"). However, the conclusions drawn from the present studies may be extended to other levels of abstraction, with the possible exception of labels for individuals, as proper names arguably do not denote categories in the same way as common names.

Although in a given context, a label often denotes a particular entity (e.g., the deictic "Please sit on the chair"), the bare label itself—"chair"—denotes a range of stimuli with the association between the label and stimulus mediated by the stimulus's typicality. This is a critical point. Although the label "chair" is discrete in that one either calls an object a chair or does not, it is more strongly associated with some chairs than others as evidenced by the relative ease of labeling typical compared to atypical items (Rosch, 1978). Overall, objects judged to be typical members of a category are more easily accessed than atypical objects (Rosch, 1973; Kail & Nippold, 1984), a finding that is easily extended to naming because most measure of access are measures of naming latency (Rosch, 1978).

A label thus involves a many-to-one mapping between stimuli (entities in the environment) and response (the name of the category)[1]. The central question of this thesis is: how does learning and using these names affect the representations of objects denoted by the names.

1.2 Historical Origins

The notion that recognizing an object as *something* requires more than passive perception stems from the centuries-old idea that a stimulus is inadequate to give a complete account of the resultant experience, as discussed, for instance, by Locke's doctrine of primary and secondary qualities of objects. Locke and other associationist philosophers theorized that a key component in the process of coming to recognize a stimulus is the learning of associations, producing "stimulus enrichment." James Mill in particular, suggested that experiences are the results of the compounding of innumerable "simple, complex, and duplex ideas." Making sense out of sensory experiences can thus be seen as "joining together and supplementing an elemental sensory input"

[1] This mapping is, of course, bi-directional. Language production involves activating appropriate labels given some concept. Language comprehension involves activating appropriate concepts given a label.

(Tighe & Tighe, 1966 p. 354). This core assumption that stimuli are somehow enriched through experience was further elaborated by Tichener and Helmholtz, but perhaps the most notable elaboration came in the form of William James's law of "dissociation by varying concomitants." James theorized that "What is associated now with one thing and now with another tends to become dissociated from either." James admitted that the nature of this process "is a little bit of a mystery" but gave two mechanisms by which such experience with associations may result in improvement in sensory discrimination. The first, most relevant to the present discussion, was what later came to be known as the doctrine of acquired distinctiveness[2]: "The terms whose difference comes to be felt construct disparate associates and help to drag them apart" (1890 p. 510). Associating initially similar percepts with more distinctive percepts or responses, makes the original percepts more distinctive[3]. James's formulation took the form of a thought experiment. How does one learn to distinguish a claret from a burgundy? James writes that probably they have been drunk on different occasions and settings, and the next time we drink the wine, "a dim reminder of all those things" is recalled. "When this setting has come to each, [our] discrimination between the two flavors solid and stable." This inclusion of context as an important variable mirrored the thinking of early associationist philosophers. What James writes next is of critical importance to the present thesis:

---

[2] The second mechanism James proposed, is summed up by the quote: "a difference reminds us of larger differences of the same sort, and these help us to notice it" (1890, p. 510). An example may be of a subtle difference in color between two objects being noticed due to previously noticing of larger color difference and the resultant sharpening of discrimination along the color dimension. Interestingly, while James's first mechanism evolved into the enrichment view of stimulus discrimination which dominated behaviorist research, the second mechanism became the basis of the opposing Gibsonian view of stimulus discrimination through differentiation (e.g., Gibson & Gibson, 1955).

[3] "Two things, then, B and C, indistinguishable when compared together alone, may each contract adhesions with different associates, and the compounds thus formed may, as wholes, be judged very distinct. *The effect of practice in increasing discrimination must then, in part, be due to the reinforcing effect, upon an original slight difference between the terms, of additional differences between the diverse associates which they severally affect*" (1890, p. 511).

> After a while the tables and other parts of the setting, besides the name, grow so multifarious as not to come up distinctly into consciousness; but *pari passu* with this, the adhesion of each wine with its own *name* becomes more and more inveterate, and at last each flavor suggests instantly and certainly its own name and nothing else. The names differ far more than the flavors, and help to stretch these latter farther apart. Some such process as this must go on in all our experience (p. 511).

Remarking further on the importance of category labels, James comments that although it may seem that the difference we feel between the two wines "we should feel, even though we were unable to name or otherwise identify the terms," this difference "is always concreted and made to seem *more substantial* by recognizing the terms." James illustrates the latter idea by the following example:

> I went out for instance the other day and found that the snow just fallen had a very odd look, different from the common appearance of snow. I presently called it a 'micaceous' look; and it seemed to me as if, the moment I did so, the difference grew more distinct and fixed than it was before. The other connotations of the word 'micaceous' dragged the snow farther away from ordinary snow and seemed even to aggravate the peculiar look in question. I think some such effect as this on our way of feeling a difference will be very generally admitted to follow from naming the terms between which it obtains; although I admit myself that it is difficult to show coercively that naming or otherwise identifying any given pair of hardly distinguishable terms is essential to their being felt as different at *first* (p. 512).

Thus the role of labels in James's view is to make the named experience at once more abstract, by overlooking idiosyncratic details (e.g., the shape of the wineglass) while at the same time highlighting the differences between it and contrasting categories (different wines or types of snow)—thus making more concrete the distinction between the categories.

While the idea that categorization relies on learned associations and involves abstraction and highlighting of features has remained central to accounts of categorization, the role of verbal labels in this process has received a comparatively miniscule amount of attention. .

1.3 Category Labels and Response Differentiation

James's law of "dissociation by varying concomitants" was elaborated by Miller and Dollard into the doctrine of acquired distinctiveness/similarity of cues (Miller & Dollard, 1941). Within the behaviorist paradigm, object names, like button presses, were seen as motor responses. Consequently:

"learning to respond with highly distinctive names to similar stimulus situations should tend to lessen the generalization of other responses from one of these situations to another since the stimuli produced by responding with the distinctive name will tend to increase the differences in the stimulus patterns of the two situations[4]. (Miller, 1948 p. 174 cited in Rossman & Goss, 1951)

Rather than the stimulus representation being altered by the label as is proposed by the language-feedback hypothesis developed here (Section 2), within the framework of Miller and Dollard (1941), labels were seen to enrich and mediate the stimulus. For instance, in Goss's (e.g., 1961) formulation of the verbal mediation theory, if stimuli A and B are mediated by names C and D respectively, then it should be easier to learn to associate A and B with different motor responses assuming that the names C and D are more distinctive than the stimuli A and B. In the 1950s and 1960s, a number of researchers tested this theory with mixed results.

[4] A parallel statement was also provided on the acquired similarity of cues through common names.

In a typical experiment, participants learn to associate novel shapes that are similar to each other with easily discriminable responses such the letters A and B or nonsense syllables—typically one response per stimulus is learned—commonly called a predifferentiation or "verbal training" phase. Discrimination ability is then assessed and compared to the ability in a group that did not receive any training or received training in associating all stimuli with a common response (Arnoult, 1953). A number of studies (Battig, 1956; Lawrence, 1949; Lawrence, 1950; Rossman & Goss, 1951) found increased discriminability of visual stimuli in adults after learning to associate diverse verbal or motor associations with the stimuli, but there is no conclusive evidence that the mechanism involved is that described by Miller and Dollard. In most studies "verbal pretraining" was confounded with perceptual practice, i.e., the verbal pretraining groups received discrimination training in the form of learning verbal names for the stimuli, while the control groups not only did not learn the names, but received no training at all, only incidental exposure which did not involve practice in discriminating the stimuli. It is therefore unclear whether the increased discriminability arose from the learned associations or from perceptual practice, as theorized by Gibson and Gibson (1955). To test whether learning labels offered any benefit beyond familiarizing the participants with the stimuli—the crucial variable discussed by Gibson—Robinson (1955), a student of the Gibsons, had participants discriminate pictures of fingerprints, either pre-training them to associate different "gangster nicknames" with each print, to categorize the prints into groups of cops and robbers, or simply to produce same-different responses comparing a current print to the previous one. All three groups received equal exposure to the stimuli. The three training groups had better performance on a subsequent discrimination task than a fourth group which received no training. But neither of the verbal conditions resulted in better performance than nonverbal discrimination training. Robinson concluded that learning of the arbitrary names for the prints did not produce change in stimulus discriminability beyond that produced by simple discrimination training. Though Robinson obtained negative findings, Goss (1953) controlling for stimulus familiarization, found that learned responses significantly facilitated discrimination, but in a motor rather than perceptual domain.

Further evidence against Miller and Dollard's hypothesis that associating distinct responses makes the original stimuli more distinct through a process of enrichment, comes from the findings of Hake and Erikson (1956) who tested the prediction that positive transfer should increase with the number of responses learned for a given number of stimuli because more response-produced stimulation is made available to distinguish the stimuli (Tighe & Tighe, 1955). In accord with Robinson's (1955) finding, Hake and Erikson found no effect of label specificity on subsequent recognition of complex visual forms, concluding that "the perceptual gain resulting from labeling practice appears to occur as long as subjects have a decision to make about the stimuli on each trial (p. 167). DeRivera (1959) further confirmed this finding.

Miller and Dollard's doctrine of acquired distinctiveness did not go unnoticed in developmental psychology, see 2.3.2. Despite strong interest in the 1950s and 1960s in the role of labels in concept formation and perceptual learning (further reviewed in Section 3.2), the research has suffered from weak theoretical guidance (Katz & Zigler, 1969). Missing from these studies was the idea that actual labels refer to categories, and a learned label is a learned category. To learn a label successfully requires associating the label with category-relevant features, and largely dissociating it from category-irrelevant features. The learning *and use* of a label may therefore strengthen the category representation, but only insofar as category members have something in common. From this perspective, it is not surprising that labels may have effects only when associated with *several* items that form a category, but not when arbitrary labels are associated with individual arbitrary stimuli, as was common in the experiments conducted within the framework of acquired distinctiveness of cues.

1.4 Category Labels and Perceptual Learning

Perceptual learning within Miller and Dollard's (1941) doctrine of acquired distinctiveness was formulated in terms of stimulus and response. With the demise of the behaviorist paradigm, interest in the topic of perceptual learning waned despite an enormous amount of evidence that learning and categorization shape low-level perceptual processing (Gibson, 1953). Categorical effects on perception became viewed as evidence of domain-specific and perhaps innate processes (e.g., Liberman, Cooper, Shankwei, & Tuddert, 1967).

In the 1980s, Stevan Harnad's *Categorical Perception* (1987) renewed interest in perceptual learning. Arguing that categorical perception is pervasive in numerous domains and is therefore far from being special to phonemes, Harnad reframed the focus from cue-association to representational change. The subsequent studies of Rob Goldstone greatly elucidated the ways in which category-learning shapes perception (Goldstone et al., 2004 for review). By using parametrically varying stimuli, Goldstone was able to precisely quantify the degree to which learning to categorize things apart or together makes them more different or similar respectively.

Although disconfounding to some degree the effects of categorization from those of mere exposure, Goldstone's studies did not aim to dissociate the role of category names from that of differential responses. Participants typically learned to respond to some items as belonging to "Category A" and others to "Category B" (Goldstone, 1994) or to discriminate between individuals belonging to a "club" to those not belonging (Goldstone, Steyvers, & Rogosky, 2003). Following this training, discrimination ability was assessed. Because responses were confounded with labels, it is unknown what work, if any was done by the category names themselves. While learning to categorize even perceptually complex stimuli is far from a uniquely human ability (see Herrnstein, 1990 for an example of categorization by pigeons), habitually learning names for those categories *is* uniquely human. Considering the prominent place that language plays in human behavior, understanding how naming shapes cognitive and perceptual processing remains a key issue.

Contemporary studies, such as Goldstone's, interact in an interesting way with earlier efforts to study perceptual learning. For instance, Donderi, Seal, and Covit (1973) trained participants to associate nonsense geometric shapes with either similar or different associates. Participants' discrimination of the two shapes was tested by having them provide same/different responses under brief exposures prior to training, and a week subsequent to it.[5] Participants trained with dissimilar associates improved their discrimination more than participants trained with similar associates. The authors concluded that learned associations "affect discrimination directly, by adding another part to the mental organization …which is activated following stimulation by each discriminandum" (p. 400) (cf. Sloutsky & Fisher, 2004a). This contrasts sharply with Goldstone's view of what happens during perceptual learning. According to Goldstone, perceptual learning is about long-lasting changes to the perceptual system (1998), which may come about through a process of reorganizing existing feature detectors as well as creating new ones (Goldstone, Steyvers, & Larimer, 1996; Goldstone et al., 2004) and sensitization in regions of perceptual space that signal differences in desired responses (typically near the category boundary) (Goldstone, 1994). Goldstone further theorizes that in addition to categorization being

[5] See (Gauthier, James, Curby, & Tarr, 2003) for a very similar experiment using conceptual associates. Participants had improved visual discriminability when they were trained to associate visual stimuli of the kind used here in Section 1.5 with more distinct adjectives sets (fast, flexible, angry vs. cold, rare, sweet) compared to overlapping adjective sets (loud, nocturnal, strong vs. loud, heavy, strong).

based on featural and dimensional descriptions of objects, the categorization process partially forms the descriptions that are used (1994 p. 199). So, rather than categorization being based on bottom-up processing of features provided by the perceptual system, categorization can influence the features that are used. While this can be viewed as a more sophisticated version of the enrichment view of perceptual learning, such an account may be more compatible with the feedback framework proposed here (see Section 3) and hinted at by Goldstone (1994).

In discussing their results, Donderi and colleagues (1973) invoked Vernon's (1955) theory of schemata: "When the stimulus is weak, the power of the schema to determine perception is strong (Donderi et al., p. 400).[6] While the relevance of schema to their results is not obvious, in the framework of the theory proposed in this thesis, the schema becomes the label itself. It is not that the label adds another feature to the referent (cf. Sloutsky & Fisher, 2004a), but rather the label schematizes the stimulus by placing it into the larger context of the category, allowing its representation to be "informed" by other members of the category.

---

[6] Another demonstration of the role that categories play in perception comes from a curious study by Donderi and Kane (1965) using perceptual disappearance. When subjects are exposed to simple luminous figures in a dark room, partial stabilization of these figures on the retina reduces the effective stimulus (Mckinney, 1963). The authors trained participants to group together 2 of 3 shapes differing on a some dimension, e.g., learning to respond with "A" to a small and medium circle, and "B" to a large circle. Following this training, it was found that the stimuli grouped together disappeared more often than those not grouped together. In this paradigm segments of the display disappear in "meaningful units" and common elements disappear together (McKinney, 1963). Therefore the greater disappearance of shapes previously grouped together offers evidence that discriminably different stimuli react as they became more identical following common response training. While this is an impressive finding, the perceptual disappearance paradigm depends on subjective responses—there is no way to verify the subject's claim that a particular segment of the display actually disappeared. It is also impossible to rule out response bias from these results.

2 The language-categorization-perception link

2.1 *The Whorf Hypothesis Revisited*

The central question that Benjamin Lee Whorf asked in his writing is summarized thusly: "Are our own concepts of time, space, and matter given in substantially the same form by experience to all men, or are they in part conditioned by the structure of particular languages?" (Whorf, 1939/1956, p . 138). The answer to which he arrived, largely through introspection and informal cross-linguistic analysis is summarized by Whorf's most famous passage:

> We dissect nature along lines laid down by our native languages. The categories and types that we isolate from the world of phenomena we do not find there because they stare every observer in the face; on the contrary, the world is presented in a kaleidoscopic flux of impressions which has to be organized by our minds—and this means largely by the linguistic systems in our minds …. Language is not simply a reporting device for experience, but a defining framework of it. (Whorf, 1956 p. 213).

I will not delve into the long history of the Whorfian hypothesis. Stimulating discussions can be found in the following papers: (Boroditsky, 2003; Carruthers, 2002; Gentner & Goldin-Meadow, 2003; Gumperz & Levinson, 1996; Lucy, 1992; Slobin, 1996). Importantly, the strong version of the Whorfian hypothesis—that language determines what one can think—was never seriously entertained by any experimentalist despite enjoying wide popular appeal. Indeed, the human ability to learn additional languages is obvious proof that language is not a "prisonhouse"—a point of which Whorf was well aware.

Of experimental interest is the weaker version of the hypothesis—that people speaking different languages organize information differently—and the question more central to the present thesis, that Language (i.e., any human language) shapes information processing outside of communicative tasks. Versions of this weaker hypothesis have been receiving a great deal of support. Studies have shown effects of language on categorization (e.g., Davidoff, 2001; Diesendruck, 2003; Lupyan, Rakison, & McClelland, 2007; Nazzi & Gopnik, 2001; Yoshida & Smith, 2005), visual-discrimination (Katz, 1963; Ozgen & Davies, 2002; Gauthier et al., 2003; Pilling, Wiggett, Ozgen, & Davies, 2003; Yoshida & Smith, 2005), spatial-cognition (Hermer-Vazquez, Spelke, & Katsnelson, 1999; Levinson, Kita, Haun, & Rasch, 2002; cf. Li & Gleitman, 2002), numerical abilities (Spelke & Tsivkin, 2001; Lemer, Dehaene, Spelke, & Cohen, 2003; Pica, Lemer, Izard, & Dehaene, 2004), and representations of time (Boroditsky, 2000; Boroditsky, Ham, & Ramscar, 2002). Experiments revisiting the basis of color categories (Berlin & Kay, 1969), have provided strong evidence of linguistic effects in both color memory and perception (Daoutis, Pilling, & Davies, 2006; Davidoff, Davies, & Roberson, 1999; Drivonikou et al., 2007; Gilbert, Regier, Kay, & Ivry, 2006; Winawer et al., 2007; Ozgen, 2004 for review).

Some of these studies have provided indirect evidence for James's idea that differences between categories are made more substantial by category names. For instance in the domain of color perception, some authors (Ozgen, 2004) have speculated that color naming provided categorization practice which, over time, pulled apart differently named items and collapsed distinctions among items sharing a name—the view of labels as increasing categorization practice that through representational change produces categorical perception (e.g., Goldstone, 1994; Goldstone, Lippa, & Shiffrin, 2001).

2.2 A Paradox? Long Term versus On-line effects of verbal labels

Despite the preponderance of findings supporting the link between language and "non-linguistic" cognitive processes, there is at present no general theory about the mechanisms that give rise to these effects. The mechanisms by which language affects cognitive processing and the relevant variables involved are largely unknown. Consider the following paradox which not only has not been resolved, but has been barely addressed. There are now several studies showing that color processing is affected by linguistic terms such that individuals who speak a language that splits some part of the color spectrum into separate categories (e.g., blue and green) are more sensitive to the differences along the linguistically-defined boundary than individuals speaking a language that uses a single term ("grue") to refer to both categories (Daoutis et al., 2006; Roberson, Davidoff, Davies, & Shapiro, 2005; Winawer et al., 2007). Learning words such as "blue" and "green" may therefore provide categorization practice that results in the gradual "drawing apart" of the parts of the color spectrum with different names—the same mechanism proposed for perceptual learning (Gibson & Gibson, 1955; Goldstone, 1994; Goldstone, 1998)—and the currently dominant explanatory mechanism of the linguistic effects on color categories (Ozgen & Davies, 2002; Ozgen, 2004). If true, then once the work of the labels was done, it should not matter whether language was available on-line during the categorization task. However, when participants are placed under conditions of verbal interference, linguistic effects on color memory and perception disappear (Roberson & Davidoff, 2000; Gilbert et al., 2006; c.f. Pilling et al., 2003).

This linguistic-bleaching effect of verbal interference is also seen in other domains such as memory for events of different tenses. English and Indonesian-speaking monolinguals show memory patterns consistent with their language: better memory for different tenses in English, but not Indonesian which does not require morphological tense markers. This difference disappears under verbal interference (Boroditsky, 2003). Similarly, disruptions of language in aphasia result in nonlinguistic impairments on categorization tasks that are seemingly caused by the language loss (Section 2.5). Assuming these cognitive impairments are caused by the linguistic impairments, this could only happen if the effects of language on non-linguistic cognition have an on-line component rather than just the product of increased practice in making certain types of category judgments.

Given that language, perception, and higher-level conceptual processes continuously interact (Rumelhart & McClelland, 1982; McClelland & Elman, 1986; Spivey-Knowlton, Trueswell, & Tanenhaus, 1993; Spivey, Tyler, Eberhard, & Tanenhaus, 2001; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995; Christiansen & Chater, 1999), we may expect that language should influence representations on-line. This is not surprising when we consider that human categorization is at its core a flexible and context-sensitive process requiring us to stress different features of a stimulus as we categorize at different levels of specificity (a toothbrush or *my* toothbrush), and given different task demands (e.g., sorting items by common occurrence vs. common properties). While considerable progress has been made in understanding how context can affect which features are stressed in distributed representations (Rogers & McClelland, 2004), the role that language plays in this process of representational shift—either perceptual or of higher-level representations—has not been investigated.

The experiments detailed in this work are limited to exploring the effects of labels on visual representations (but see Sections 2.6.1 and 2.6.2). I will claim that the effects of category labels are best understood as a process of neural feedback. Supporting neural mechanisms of feedback/reentrant visual processing are discussed in Section 3.3.2. Central to the thesis is an exploration of long-term versus on-line components of labels. The studies presented in Section 4 will aim to show that in addition to any deliberate and serial processes engaged by naming, processing labels has an effect on multiple objects in parallel. Because labels and categorization are inextricably tied, investigations of verbal labels necessarily include discussions of conceptual effects, for instance, the role of conceptual categories in visual processing.

Although I will focus on visual processing, the resulting framework can shed light on a variety of Whorfian effects that have been recently reported and on the cognitive consequences of language-acquisition by children (Balaban & Waxman, 1997; Colunga & Smith, 2005; Xu, 2002; Gentner & Loewenstein, 2002; Sloutsky & Fisher, 2004a; Waxman & Markow, 1995; Waxman, 1999; Yoshida & Smith, 2005). An overview of effects of language on categorization in development follows.

2.3 Evidence from Development

2.3.1 Labels and Categories in Infancy

Much of what we know about the role of language in cognitive processing comes from developmental studies that have investigated the consequences of language acquisition, particularly on categorization and concept learning. This section is not intended to be a full review of the voluminous literature on the effects of labels on infant cognition. Instead, I briefly review a few of the approaches taken to study how labels may affect category learning and category representations, and the major conclusions of this work.

The most basic findings is that infants care about words. Ten to fourteen month-old infants devote more attention to objects that had been labeled with a noun, than to objects that had not been labeled (e.g., Baldwin & Markman, 1989). In a task designed to test the ability of 9-

month-olds to individuate objects (a ball from a duck), Xu (2002) reported that providing labels for the objects improved infants' ability to conceive of the two objects as separate and has framed these results in a theory of labels as "conceptual placeholders." Waxman (e.g., Waxman & Booth, 2003; Waxman, 2004) has also implicated word learning in conceptual organization. For instance, Waxman and Hall (1993) showed that labels direct the attention of infants to taxonomic categories in a task in which infants choose between a thematic and a taxonomic alternative. These researchers have also shown how different types of words, for instance, nouns versus adjectives, lead infants and older children to form different types of categories (e.g., object-based versus property-based), (Waxman & Booth, 2003) and have argued that words serve as "invitations for categories" (Waxman & Markow, 1995). Waxman and colleagues conclude that at the age of about 12 months, infants assume that words refer to object categories (Waxman & Hall, 1993; Waxman & Markow, 1995).

Studies investigating effects of labels on infant categorization typically show that the presence of an often familiar label (e.g., "Look! A duck!") during the familiarization phase produces a novelty preference for out-of-category items in a subsequent testing condition. If a novelty preference is not shown in a no-label condition ("Look at this one!"), this difference across conditions is taken as evidence that labels facilitate categorization (Balaban & Waxman, 1997; Fulkerson & Haaf, 2003; Waxman & Markow, 1995; Waxman & Booth, 2003). However, because most of these studies use familiar objects (e.g., dogs, cars, horses), it is not clear whether labels resulted in the teaching of the category, or a reactivation of a previously learned category (Oakes & Madole, 2000b; Plunkett, Hu, & Cohen, 2007). Considering that even very young infants (4-5 month-olds) show out-of-category preferences (as measured by longer looking times when familiarized with chairs and tested on couches and beds) without any assistance from labels (Behl-Chadha, 1996), Waxman and colleagues' findings are more compatible with the interpretation that labels "reactivated" existing categories rather than taught infants new ones. In a study in which novel objects were used (Booth & Waxman, 2002), labeling effects on 14-month-olds were observed only when infants were provided with information about the function of the objects (which had the effect of making the objects more familiar). A recent study by Fulkerson and Haaf (2006) used stimuli argued to be novel to 12-month olds (the letters H, X, T, and G), and showed that labeling affected 12-month-old infants' categorization of these objects. Yet, because the labels were presented both during familiarization (e.g., "look a fep!") and during test ("can you show me the fep?"), it is again impossible to know to what degree the labels are participating in changing the representational structure of the categories being learned.

One way to tease apart category learning versus category reactivation is to use novel items in a "process-oriented" paradigm (Oakes & Madole, 2000a). A well-established methodology for studying the formation of categories in infants is the design pioneered by Younger and colleagues (e.g., Younger & Cohen, 1985; Younger, 1990). Testing the hypothesis that infants can form categories comprised of correlated attributes, Younger (1985) familiarized infants with pictures of cartoon animals constructed to vary on a number of dimensions such as tail thickness and leg length. Following familiarization, infants are presented with paired test-trials and proportion of time spent looking to each stimulus is recorded. The degree to which infants have formed a categorical representation of the stimuli can be judged by measuring infants' familiarity with a stimulus typical of the hypothesized category (e.g., a prototype stimulus containing average feature values). For instance, if one category comprises thick-tailed, long-necked animals ($categ_1$), and another of thin-tailed, short-necked animals ($categ_2$), then showing a preference for the boundary novel animal that is medium-tailed and medium-necked over a novel stimulus typical of $categ_1$ or $categ_2$, demonstrates that infants formed two categories. Showing the reverse pattern—finding the average stimulus more familiar—indicates that infants formed a single broad category. Altering the degree of co-variation of properties among the familiarization stimuli resulted in the formation of either one broad, or two narrow categories in 10-month-old infants (Younger & Cohen, 1985). This methodology is very well suited for exploring the role of labels in category formation of infants. By introducing labels into the familiarization phase and comparing the looking patterns resulting from different labeling conditions, one can deduce how labels augment the learned category representations. Plunkett, Hu, and Cohen (2007) did just that, replicating Younger's original design, but adding novel labels ("dax" and "rif"). In the label conditions, the labels either reinforced the perceptually defined category boundaries, were paired randomly with the stimuli, or comprised a common label presented with each stimulus. Crucially, these labels were presented *only* during the familiarization phase of the study. The authors found that the common-label condition results in the formation of one broad category rather than two narrow categories, while presenting random labels prevented the infants from forming any category at all. Because the familiarization stimuli formed two separate perceptual clusters (long-necked short-legged animals, and short-necked long-legged animals), they could be categorized successfully with or without labels. Yet, when presented with a common label, the category boundary collapsed. Categorization (i.e., sensitivity to a prototype) was disrupted entirely when infants presented with labels that were inconsistent with any category structure. The authors thus concluded that labels can override (at least some) perceptually defined categories in infants as young as 10 months.

This result (as well as the numerous studies mentioned in Section 2.1) conflicts with the "cognitive priority" hypothesis that arguably still dominates the field. According to this hypothesis (also known as the "cognition hypothesis") language reflects but does not shape concepts (Cromer, 1974; Johnston, 1985; Mccune-Nicolich, 1981; Nelson, 1974; Slobin, 1985). For instance, in the domain of spatial concepts (e.g., learning to correctly categorize "above" or "in" object relations), proponents of this view posit that children's cognitive grasp of a particular spatial concept is the main determining factors in the acquisition of the corresponding spatial term. For instance, Mandler (1996) has proposed that children select linguistically needed spatial distinctions from an existing array of spatial concepts formed pre-verbally and several theories concerning acquisition of spatial terms suggest that young children map the spatial terms of their language directly onto their nonlinguistic spatial knowledge (Clark, 1973; Johnston, 1981; Johnston, 1988). Despite their considerable differences, what all these accounts have in common is the idea that the meanings expressed in children's early spatial terms are formed independently of language (see Gleitman & Papafragou, 2005 for an extended discussion of this view).

A number of researchers investigating the acquisition and organization of concepts in infants and children have stressed the importance of linguistic input for relational concepts (e.g., spatial concepts) (Gentner, 1982; Schlesinger, 1977; Vygotsky, 1962; Xu, 1999). For example, Bowerman (1989; 1996) asserts that young children's concepts of space interact with language in the acquisition of semantic spatial categories. Although she believes that spatial perception and cognition are important prerequisites in the acquisition of language-specific semantic categories, she argues that nonlinguistic cognition alone cannot account for the array of meanings expressed in language-specific semantic categories. Rather, nonlinguistic spatial cognition and linguistic input continually interact in the formation of language-specific semantic categories, although the nature of this interaction changes as children's spatial cognition and linguistic knowledge develop (Bowerman & Choi, 2001). An example of this type of interaction is seen in studies that demonstrate developmental trajectories as a function of language. For instance, McDonough, Choi, and Mandler (2003) familiarized Korean- and English-learning infants of 9, 11, and 14 months to tight- versus loose-fit containment events, a spatial distinction that is linguistically relevant for Korean but not English speakers. When viewing new examples of both types of relations, infants in each language group discriminated between the familiarized and novel relation. In contrast, only 38% of adult English speakers compared to 80% of adult Korean speakers correctly differentiated between tight- and loose-fit events. The decrease in sensitivity to the tight/loose feature in English-speaking children corresponds with increases in vocabulary level and the production of the word "in." Children who produced "in" or who had high vocabulary levels were less sensitive to the difference between tight-in and loose-in than those who did not produce "in" or had low vocabulary levels (2006). It is important to note here that language was not instrumental in establishing initial sensitivity to the dimension *tightness-of-fit* (see also Hespos & Spelke, 2004). Rather, as children begin to build their vocabularies, language plays a role in directing young children's attention to particular aspects of their environment.

2.3.2 Labels and Subjective Similarity

In addition to the quickly growing literature on effects of labels in infancy, there exist a number of studies investigating the effects that labels have on similarity relations. In a series of studies with children in the 5-9 year range, Katz (Katz, 1963; Katz & Zigler, 1969; Katz, Karp, & Yalisove, 1970) briefly trained children to associate nonsense labels with arbitrary geometric shapes, comparing children's similarity ratings of items given the same names to items given different names. As predicted, items given different names were rated as being more different, an effect mediated by the items' initial similarity—identical labels modified perceptual judgments made to dissimilar stimuli more than to similar stimuli (Katz & Zigler, 1969). Notably, effect of labels were observed only for the younger children. Katz speculated that because older children label objects automatically (Spiker, 1956) the measurable difference between the label and no-label groups disappears. More recently, Landau and Shipley (2001) provided an essentially identical demonstration with two- to three–year-old children as well as adults, showing that judgments of identity of parametrically varying stimuli are affected by the names given to object "standards." These effects of common labels making stimuli more similar bear some similarity to Sloutsky and Fisher's (2004a) findings of the effects of labels on what is essentially an subjective-similarity task (responding to questions in the form "is this one more like ___ or ____"). Unlike Katz's studies however, Sloutsky and Fisher's did not involve training.

One difficulty with interpreting findings that involve subjective similarity ratings is that it is unclear whether the labeling actually affects stimulus or category representations, or whether the altered similarity judgments reflect expectations: after being told by the experimenter that A and B have the same name, the children may well give them a high similarity rating judging that because they have the same name, they ought to be similar. Goldstone et al. (2001) note that similarity judgments are particularly prone to strategic considerations of this sort— "this is a dax, and that's a dax, so they must be similar."

On the one hand, there seems to be indisputable evidence that shared labels affect children's judgments (see also Colunga & Smith, 2005; Sloutsky & Fisher, 2004b; Sloutsky & Fisher, 2004a; Smith, Colunga, & Yoshida, 2003; Yoshida & Smith, 2005). On the other hand, it is difficult to know what exactly the labels are doing. Do they cause actual representational change or can their effects be explained by strategic biases? Researchers interested in effects of labels on *representations* should therefore exercise caution in using subjective similarity and explicit comparison tasks.

2.4 Redundant Labels Facilitate Category Acquisition

Previous research has failed to dissociate category names from category responses (1.3, 1.4), and, with few exceptions (Robinson, 1955) failed to control for familiarity effects, confounding category-learning with general familiarization with the stimulus set. Evidence from development, particularly process-oriented studies using labels and novel categories (Plunkett et al., 2007) suggest a causal involvement of labels in category formation. Nevertheless, these studies leave open the question of whether the effect of labels arises strictly from the additional information they provide or whether labels facilitate category formation even when they are informationally redundant. On the first account, labels aid category formation through two mechanisms (Bloom & Keil, 2001). First, they provide cues regarding the relative importance of some objects. Thus, a child observing adults naming a group of objects as “chairs,” provides a cue that these objects, by virtue of being named, are important and that the child should attend to them. Second, they provide additional information regarding the category structure that needs to be learned. Thus, observing the label “chair” be paired with different types of objects is a cue that these objects somehow belong together, even if the basis of their “belongness” cannot be ascertained {Brown, 1986 498 /id /ft " Chapt. 5"}.

In addition to these two mechanisms, labels may facilitate category formation even when they do not contribute extra information. Learning to associate different exemplars with a common name may result in representations that highlight category-relevant features and suppress idiosyncratic attributes that are irrelevant to category placement. This idea is discussed in more detail in Sections 2.6.2 and 3.1. Thus, the question that still remains is whether labeled categories are easier to acquire *because* they have a name—even when categorization can be performed without relying on labels.

Lupyan, Rakison, and McClelland (2007) attempted to answer this question by using a design that (1) disconfounded category names from category responses and familiarization (2) used an objective dependent variable: category-placement rather than judgments of perceived similarity which is likely to affected by strategic biases, i.e., are A_1 and A_2 judged to be similar because they are perceived as such or merely because they share a label (Goldstone et al., 2001), (3) compared the effect of verbal labels to other non-linguistic category associations.

The basic task required participants to learn to classify 16 "aliens" (Gauthier et al., 2003) into two categories—those to be approached, and those to be avoided—by responding with the appropriate direction of motion. These behavioral categories were chosen specifically because they exemplify typical categories learned by non-human animals. The category distinction involved subtle differences in the configuration of the "head" and "body" of the creatures.

Participants were randomly assigned to a *label* or *no-label* condition. On each training trial, one of the 16 aliens appeared in the center of the screen. After a 500ms delay, an outline of a character in a spacesuit (the "explorer") appeared in one of four quadrants—left, right, top, or bottom of the alien. Participants were instructed to respond with the appropriate direction key depending on the category of the alien (approach vs. escape). Auditory feedback—a buzz or a bell—sounded 200 ms after each response. In the *label* condition, a printed label then appeared next to the alien. After another 1500 ms, the stimulus was erased and a fixation cross marked the start of the next trial. The total duration of each trial was equal in the two conditions. The pairing of the label ("leebish" or "grecious") with the category (move away, move towards), and the perceptual category (family 1, family 2), was counterbalanced across participants. There were 144 training trials. All the participants received the same number of categorization trials, the only difference between was the groups was the presence or absence of category labels accompanying

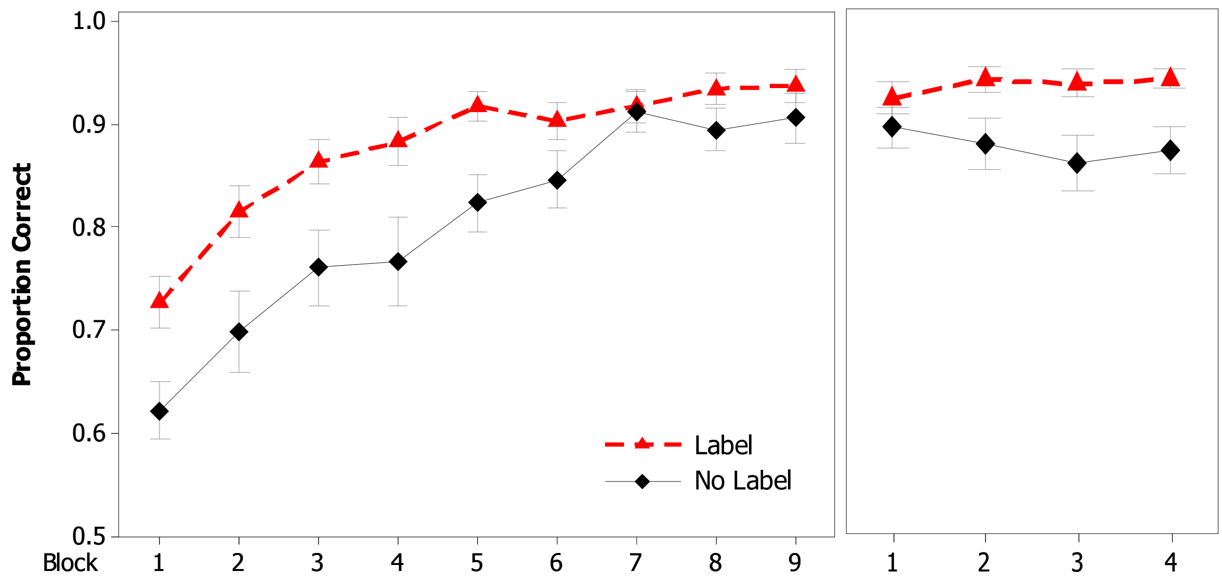

Figure 1: Mean classification accuracy for training (left) and testing (right). Learning to categorize named categories is easier than unnamed categories, and results in more robust category knowledge. Bars show one SE of the mean (from Lupyan, Rakison, and McClelland, 2007).

each response. In the *label* conditions participants learned named categories, and in the *no-label* condition they learned unnamed categories. Because the labels were perfectly correlated with the behavioral categories, they constituted entirely redundant information. After completing the category-training phase, participants' knowledge of the categories was tested without feedback or labels. Since no reinforcement or correction was provided in the testing phase, category knowledge may deteriorate over time.

Participants in the label condition performed better overall and learned the categories significantly faster than participants in the *no-label* condition (Figure 1 left). Participants who learned the categories with labels also retained their category knowledge throughout the testing phase, whereas those who learned the categories without labels displayed decreasing accuracy over time (resulting in a significant group × block interaction) (Figure 1 right). There was no difference in latencies between the two groups in any part of the experiment, suggesting that the advantage in the label condition did not arise from greater time spent on task.

To summarize, participants learning category names for novel stimuli were faster to learn correct behavioral responses for the stimuli. This was true despite all participants undergoing the same amount of experience discriminating and categorizing the stimuli. Once the categories were learned by both groups, participants in the label group showed more robust category knowledge as evidenced by continued good performance when feedback and labels were removed.

These experiments directly tested the idea that labels make category differences “more concreted” (James, 1890, p.333). Labeled categories were easier to learn even when the labels were entirely redundant, contributing no additional information. Why did the labels help? One simple way to conceptualize the effect of labels is that they pushed apart the visual representations. By associating the perceptually-similar stimuli with perceptually distinct stimuli—the labels—the representations of the aliens were “pushed apart” (Goldstone, 1994; Goldstone & Barsalou, 1998; Goldstone, 1998). But how does this actually work? Is the effect of labels a long-term “warping” of the conceptual/perceptual space, or is it an on-line affect that requires co-activation of the category names? Why do labels have an effect on category-learning over and beyond already distinct behavioral categories? If labels just produce long-term representational change—that is, do their job during learning—then it should not matter whether an object is named—the labels have already exerted their influence. Yet, it does.

2.5 Evidence from Language Loss

If language not only facilitates category-learning, but also provides on-line support for categorical thinking (Goldstein, 1948; Noppeney & Wallesch, 2000; Vignolo, 1999), then defects in language may produce impairments in categorization tasks. This idea was discussed at length by Kurt Goldstein, a German neurologist and aphasiologist (1924; 1948). Goldstein considered the ability to fixate thoughts one of the main functions of language, remarking:

> Language is not only a means to communicate thinking; it is also a means to support it, to fixate it. Defect in language may thus damage thinking. (Goldstein, 1948, p.115)

Though aphasic patients have impairments on a wide range of tasks not requiring overt language use (Lupyan, in preparation), the most consistent and profound deficits are seen in object categorization tasks. While it is difficult to isolate effects of language impairments on categorization from those of semantic impairments, this is possible by focusing on tasks in which performance is correlated with language impairments controlling for semantic impairments. The resulting tasks have something in common: most require the patient to selectively attend to particular features while “ignoring” others. Cohen, Kelter, and colleagues, concluded that aphasics have a “defect in the analytical isolation of single features of concepts.” (Kelter, Cohen, Engel, List, & Strohner, 1976; Cohen, Kelter, & Woll, 1980; Cohen, Woll, Walter, & Ehrenstein, 1981). All tested subtypes of aphasic patients are “deficient if the task requires isolation, identification, and conceptual comparison of specific individual aspects of an event,” but are equal to controls “when judgment can be based on global comparison” (Cohen et al., 1980). To illustrate, consider patient LEW first described by (Druks & Shallice, 1996; Druks & Shallice, 2000) and further tested by Roberson, Davidoff, and Braisby (1999) and Davidoff and Roberson (2004). LEW is anomic, but has excellent comprehension, and is severely impaired on taxonomic-grouping tasks with not only complex stimuli (faces), but even the simplest perceptual stimuli (e.g., colors, and simple shapes), The authors write:

> For such classifications [taxonomic groupings], the concept and the name are in effect the same thing and LEW is without names to assist the categorical solution. Where patients such as LEW can name, they can categorize. (Davidoff & Roberson , 2004, p. 166).

However, LEW is not impaired in making all types of categorization judgments. Rather, his impairments are isolated to those requiring isolation of concepts. For instance, while he can recognize which objects share a common function (a type of global comparison through common context), he is deficient in grouping together objects of common size or color—tasks requiring isolating the perceptual dimension of color or size from the rest of the object's features.

If the unvailabaility of category names is responsible for the poor performance, can providing names ameliorate the deficit? There is some evidence that it can. Koemeda-Lutz and colleagues (1987)showed that performance on aphasic patients on tests of category and property verification improves when names are provided. Based on these results the authors theorized that verbal label can help the aphasic patients to identify and keep in memory the property or category needed to make subsequent judgments, concluding that: words may function "as useful devices for keeping a clear orientation in cognitive searching processes and for focusing on specific aspects of meaning… Red cherries and red bricks may be judged to be alike mainly via what is concentrated and coined in the verbal label 'red'" (pp. 332-333).

2.6 Two Tests of On-Line Effects of Labels

2.6.1 An Effect of Verbal Interference on Categorization of Familiar Objects by Color, Size, and Thematic Relationship

If some of the conceptual "dissolution" seen in aphasics is due to the failure of language to maintain conceptual representations, then normal subjects when placed under conditions of verbal-interference may exhibit some of the same symptoms exhibited by aphasic populations: idiosyncratic and broadened semantic fields (Caramazza & Berndt, 1982; Grossman, 1978; Mccleary, 1988), "blurred group structure" (Kelter, Cohen, Engel, List, & Strohner, 1977; Zurif, Caramazza, Myerson, & Galvin, 1974), and a difficulty in isolating perceptual dimensions (Cohen et al., 1980). None of these questions have been unanswered.

The experiment described in this section tested the hypothesis that verbal interference produces in normal participants some of the same categorization impairments seen in aphasic patients. A positive finding would provide further support to the idea that impeding access to category labels can affect categorization performance, thus demonstrating the on-line interaction of language and categorization.

Recall that the most consistent categorization impairment in aphasic patients is a failure to isolate and identify specific (mostly perceptual) dimensions (Kelter et al., 1976; Cohen et al., 1980; Cohen et al., 1981). To the extent that verbal interference disrupts the on-line linguistic component responsible for selecting and comparing specific dimensions (perhaps through the use of dimensional terms like "color"), performance should be more affected in tasks requiring judgment of specific, particularly perceptual, features.

To test this hypothesis, participants performed an odd-one out categorization task in which given three objects, they had to choose one that did not belong based on color, size, or thematic relation. Based on the evidence that aphasic participants with mild or no semantic impairments have particular difficulties with tasks requiring isolation of perceptual features, it was predicted that verbal interference would have a stronger effect on color and size judgments than for more general association (thematic relation judgments). The design for this experiment was borrowed from Davidoff and Roberson's (2004) Experiment 7, conducted with the anomic patient LEW. The patient showed the expected pattern of results.

Full methods and results will be reported elsewhere (Lupyan, in prep). Participants were instructed that they would see triads of objects (grayscale pictures in the first experiment, words in the second experiment). For each trial, they were instructed to select the object that did not belong based on size, color, or theme. The criterion for the block was prominently displayed at the start of each block and at regular intervals thereafter. Examples of each type of categorization were provided and it was stressed that color and size referred to the objects the pictures or words represented. Thematic trials were composed of triads that required picking out a common relation (e.g., table / door / doorknob; watch / lock / key), common context (potato / balloon / cake) or complimentary use (chisel / hammer / nail).

Verbal interference (VI) was implemented as a within-subjects manipulation. Participants were instructed that before some trials they would see a 9-digit number and their task was to rehearse it as they performed the subsequent categorization trials. After 5 trials, they attempted to select the correct number using 4-alternative forced choice. The incorrect choices differed from the correct choice by a transposition of 2 adjacent digits in the 3-8 positions of the 9-digit number. The verbal-interference and no-interference conditions were interleaved, alternating every 5 trials. The first experiment was conducted using picture stimuli (Rossion & Pourtois, 2004). The second experiment used identical triads with words instead of pictures.

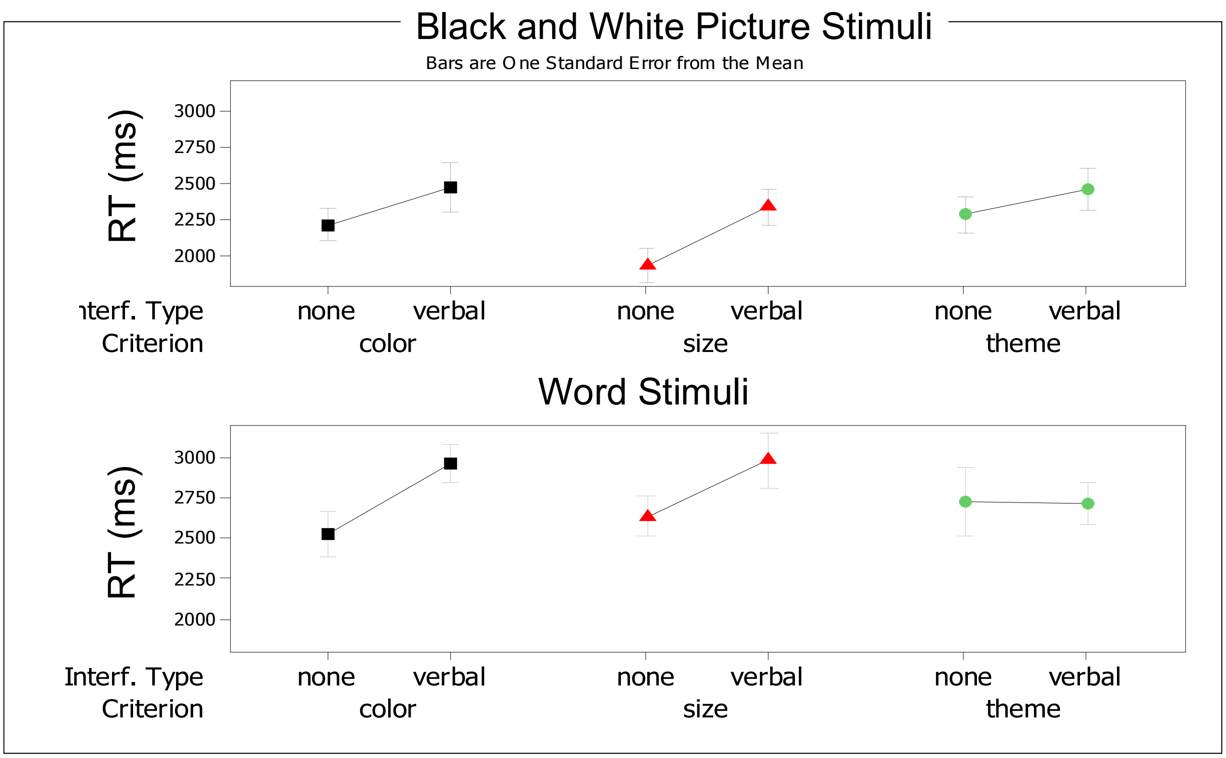

Figure 2: A comparison of baseline performance on the odd-one out task to performance under verbal interference (rehearsing number strings). Color, Size, and Theme refer to the type of response required of the participant. Response types are blocked. Error bars show between-subject SEs.

Verbal interference resulted in longer reaction times. No consistent effects on accuracy effects were found. Of primary interest is the differential effect VI had on the blocks requiring the isolation of dimensions (*color* and *size*), compared to the *theme* block which demanded broader and more holistic processing of object-relations. Results were largely consistent with the hypothesis that verbal interference impedes categorization that requires isolating particular dimensions. The effect of VI on color and size blocks was greater than on theme block (the interaction between color and theme is marginal) (Figure 2 top).

This experiment using picture stimuli leaves open the possibility that the interaction between interference and block types hinges on perceptual rather than conceptual factors. Almost all the triads in the color block used living things, while most theme-trials involved artifacts. Alternatively, the effect of VI may have been to slow down semantic access from pictures, by for instance, slowing down picture-naming. While this does not explain why color/size blocks should show a greater slow-down than the theme block, this explanation is at odds with the idea that the interaction is caused by impaired linguistic feedback that is used in isolating dimensions of stimuli.

In a second study, I attempted to replicate the effect using word stimuli instead of pictures. If the effect of VI hinges on perceptual differences among the blocks, the interference between VI and block-type should now disappear. If VI has its effect through by slowing picture-recognition (e.g., by interfering with picture-naming) semantic access from pictures, then VI should now have less or no effect on RTs. If, however, the effect is due to VI interfering with the "analytic decomposition" of categories (e.g. Cohen et al., 1980), we should see the same pattern as in Experiment 1. In fact, using words may lead to a stronger interaction because word stimuli may limit the use of any extra-linguistic information provided by the pictures. As can be seen in the bottom panel of Figure 2, using words as stimuli did not eliminate the differential effect of VI, and in fact made it stronger.

A possibility that remains is that the thematic judgments are, for some reason, just more impervious to any type of interference. An additional experiment employing a visuospatial rather than verbal interference task, in which participants were asked to remember a pattern of dots on a grid (Gilbert et al., 2006), failed to find an interaction between type of categorization (color, size, theme) and visual interference condition.

In a task requiring extracting a particular dimension such as color, and then comparing features on that dimension across several items, impeding access to verbal labels impairs performance. These results are consistent with the possibility that object and dimension labels ("watermelon", "color") interact to allow for more efficient extraction of perceptual information.

Because correct responses on the theme block do not require isolating a specific dimension, access to dimensional labels such as color is less important, and verbal interference has little effect on performance.

2.6.2 An Effect of Labels on Representations as revealed by a Recognition Memory Paradigm

Categorizing—responding in the same way to discriminable stimuli—requires abstracting over dimensions irrelevant to the category. For instance, when categorizing an object as a chair, it may be advantageous to ignore extraneous factors such as color, since chairs come in various colors. This type of abstraction becomes a liability when the task depends on faithful representation of properties both relevant and irrelevant to the category such as discriminating a previously seen item from a novel item from the same category.

The tradeoff between categorization and memory has been explored most intensely in the context of constructive memory (Bartlett, 1932) and abstraction of learned material (Bransford & Franks, 1971; Franks & Bransford, 1971; Posner & Keele, 1968). In some ways the results reported below borrow from these classic paradigms in that categorizing an item as a category member tends to involve overlooking its idiosyncratic features and encoding it more coarsely. However, the reason why overtly labeling highly familiar items is predicted to produce poorer memory for the named items is predicted to have a different explanatory mechanism.

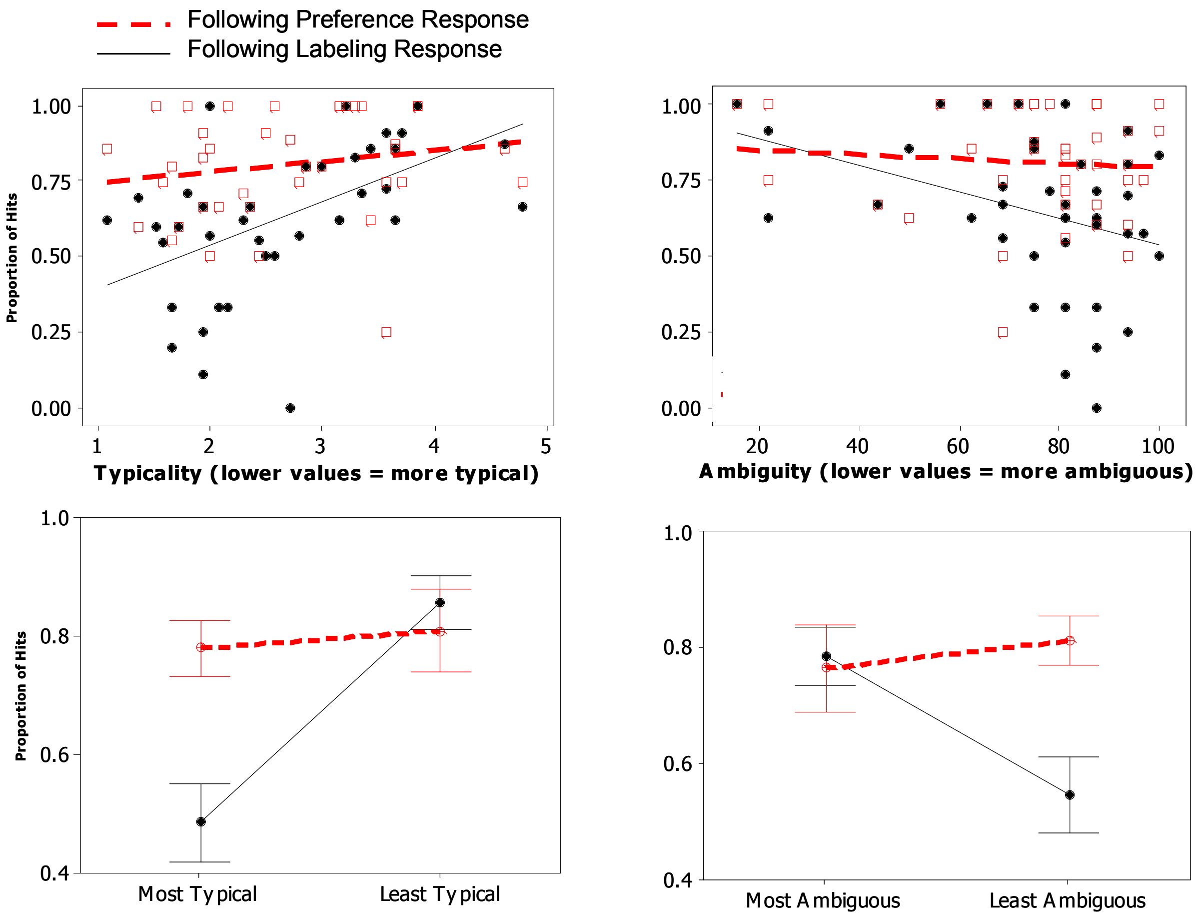

Figure 3: Recognition memory for items that were overtly classified (labeled) during the study session, or those to which a non-category related response was made (preference). Memory is shown as a function of typicality (left) and ambiguity (right). Effects of labeling are greatest for the most typical and most unambiguous items, that is, those that are most strongly associated with the label. Bars show 1 SE from the mean (between-subject).

Because labels such as "chair" denote entire categories, overtly labeling a particular chair may alter the competition between bottom-up and top-down sources of activation, resulting in a representation of a particular chair that is more influenced by previously encountered category members. This representational augmentation is predicted lead to poorer memory for specific exemplars (i.e., poor within-category memory) because successful recognition of previously encountered items depends on a match between the retrieval cue (typically the test item), and the contents of memory. The closer the match between the two, the more familiar the item seems, and the more likely it is to be correctly recognized (McClelland & Chappell, 1998; Shiffrin &

Steyvers, 1997). By distorting the original items representations, the representation of the item presented during the familiarization phase is more likely to mismatch the item presented in the test phase.

Consistent with these predictions, Lupyan (2007) found that in a recognition memory task, adults had poorer memory for pictures of chairs, tables, and lamps which they explicitly labeled as “chairs,” “tables”, and “lamps” compared to images they did not explicitly label, responding instead with preference judgments (i.e., whether they liked the item or not). Making preference judgments in this speeded task was predicted to partially block the largely automatic labeling response. The effect of labeling the pictures was reflected in the subsequent hit rates during the test phase of the experiment. Participants had dramatically lower hit rates for labeled items than to items to which they made a response unrelated to the object’s category—even though they were told to try to remember each item the best they could. The effect of labeling was strongly mediated by typicality: recognition was most affected for the most typical and unambiguous category members, as would be predicted if the effect of labeling depends on the strength of the association between the label and presented exemplar (Figure 3).

Subsequent studies ruled out explanations based on task-demands and accounts based on depth-of-processing accounts, showing for instance, that the effect persisted even when participants did not know at the time of seeing each image in the study phase, what kind of response they would be asked to make. A subsequent experiment showed that in a typicality-judgment task, simply including a label (how typical is this chair) resulted in more typical judgments than when the prompt omitted the label (how typical is this object), thus confirming the prediction that labels make objects more typical (this effect could not be explained by some items having their category membership disambiguated by the label).

These results suggest that the act of labeling a familiar stimulus with a category label augments its representation, making its later recognition more difficult due to a mismatch between the encoded representation and the item shown at test. Interestingly, the items most affected by being overtly labeled were the most typical exemplars, because it is these items that are most strongly associated with the label. As a result, the top-down feedback evoked by the labeling response appears to affect these items more than atypical and ambiguous items, which, although have more potential to be incorporated into the broader category, have only a weak association with the label and so their representations appear to be unaffected by the labeling response. This is the cornerstone of the Language Feedback Hypothesis discussed in the next section.

The idea that categorization and within-category memory are inversely related also leads to the prediction that failure to categorize may lead to superior within-category memory. Indeed, Heaton and colleagues found that while low-functioning autistic children were impaired in their ability to make category-appropriate color choices in response to a given color, they actually exceed the performance of normally developing children in a subsequent within-category memory test (Heaton, Ludlow, & Roberson, 2007).

3 The Label Feedback Hypothesis (LFH)

3.1 A Conceptual Overview

The primary claim of the LFH is that because labels most reliably co-occur with category-typical features, the use of a label distorts "lower-level" (here, visual) representations in the direction of category-typical values. This distortion is produced by a process of feedback. One implementation of this feedback process is presented in Section 3.2 and a discussion of neural mechanisms that may be involved is presented in Section 3.3.2.

The process by which labels modulate lower-level representations relies on the following assumptions: (1) Labels denote categories.[7] (2) *Within a particular context* categories have distinctive and diagnostic features.[8] (3) Labels are most strongly associated with the typical/diagnostic features because it is they that co-occur most reliably with the label. In such a framework, as the label (e.g., some form of lexical representation) becomes active, its activity feeds back onto both conceptual and more perceptual representations, altering the representation of current object or objects, highlighting their category-typical aspects while reducing their

---

[7] The use of a label requires placing the named object into a category because it is the category that is actually named. As stated by Goldstein, when we speak of "table" we do not "mean a special given table with all the accidental properties, but we mean 'table' in general. We employ the word "table" in this categorical sense when naming a particular table." Goldstein, 1936; as cited in Noppeney an Wallesch, 2002 p. 376).

[8] This is not an argument for categories being defined by necessary and sufficient features. Rather, in a particular context, certain features are more relevant than other features. In a task that involves categorizing objects by their color, color features are most relevant. In a task involving object categorization , for instance, classifying objects as either chairs or tables, shape rather than color becomes relevant and features like *has_armrests* become diagnostic. So while not all chairs have armrests, their presence is unambiguous evidence for the presence of a chair. While the presence of armrests speeds up classification of an object as a chair in the context of tables (Lupyan, 2007a). Though the relevant features differ depending on context, some, like the trunk of an elephant, are diagnostic of the animal's category in a wide range of contexts.

idiosyncratic aspects, thus making the objects more typical of the named category, and in effect, more coherent objects.

In its simplest form, the language-feedback hypothesis predicts that learning labeled novel categories should be easier than learning unlabeled categories even when the labels provide no additional information—the prediction was confirmed by the experiment presented in Section 2.4. The LFH explains this finding in the following way: because as labels co-occur with category exemplars, the associations between the label and various features varies as a function of their diagnostic value. Features irrelevant to the category distinction become dissociated from the label. Features that predict category membership (and thus the label), become more strongly associated with it. The presence of the label thus allows the diagnostic features to build up excitatory links to each other which results in more categorical representations, that is, representations with greater between-category distances and smaller within-category distances.

The LFH also predicts that the category-learning advantage of labels comes at the cost of accurate representations of individual objects. If the naming process alters object representations, people should exhibit poorer memory for objects that they name compared to objects for which the naming process is interrupted by, for instance, having them make a response unrelated to the category. Indeed, remembering particular exemplars (e.g., a particular chair) is more difficult when it is named insofar as the distortion caused by category labels creates a study-to-test mismatch. This prediction is supported by experiments discussed in section 2.6.2.

3.2 A Simple Implementation

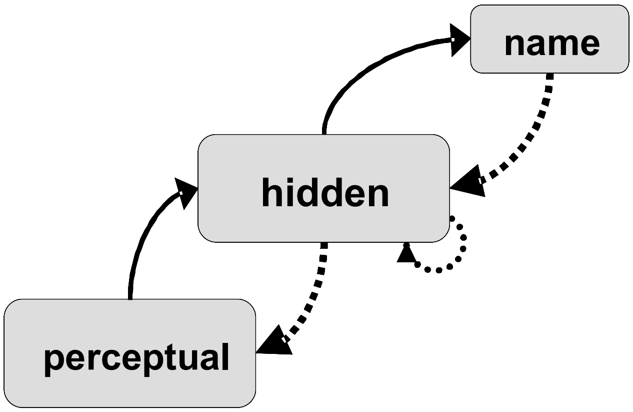

Figure 4: A schematic neural network architecture for exploring on-line effects of labels on item representations. See text for description.

A schematic of a simple implementation of label-feedback model is shown in Figure 4. Here, the model is implemented as a fully-recurrent neural network (Rumelhart, McClelland, & the PDP Research Group, 1986). Solid lines denote feedforward connections and dashed lines denote feedback connections. In this implementation, the perceptual layer is provided with a feature-based input of a current object (e.g., a given chair). During training, the model learns to produce names (i.e., produce the label "chair" given a chair and "table" given a table), and comprehend names (given the label "chair," it activates properties characteristic of chairs)[9]. In the course of experience with labeling various objects, the label becomes most strongly associated with features most commonly associated with the category (e.g., the feature *has-armrests* for chairs), and dissociated from features not reliably associated with the category (e.g., *is-brown* for chairs). Following training, we can examine what happens to representations of category exemplars when the label is allowed to influence the representations on-line. Figure 5 shows a principal-components analysis (PCA) of the *perceptual* representations of exemplars from two categories learned in the context of labels. In Figure 5A the labels are prevented from affecting the representations on-line (i.e., the name-to-hidden connections are deactivated). The category separation is due entirely from bottom-up

[9] Given only a category label, the network, of course, cannot know which particular object is being referred to. This is precisely the point. Over time, the label becomes associated most strongly with typical and diagnostic features, a representation that can be roughly equated to a prototype of the category.

perceptual differences between the two categories. In Figure 5B, the network produces the labels itself, and the produced labels are allowed to feedback to affect the visual representations. In Figure 5C the labels are provided to the network (equivalent to a task in which an object is presented and named by the experimenter). Clear category separation is observed. The PCA analyses demonstrate that when category labels are allowed to affect the exemplar representations, categorization is facilitated (items cluster more neatly into their categories), but at the expense of accurate representation of individual exemplars (i.e., poorer recognition memory) (Lupyan, 2007).

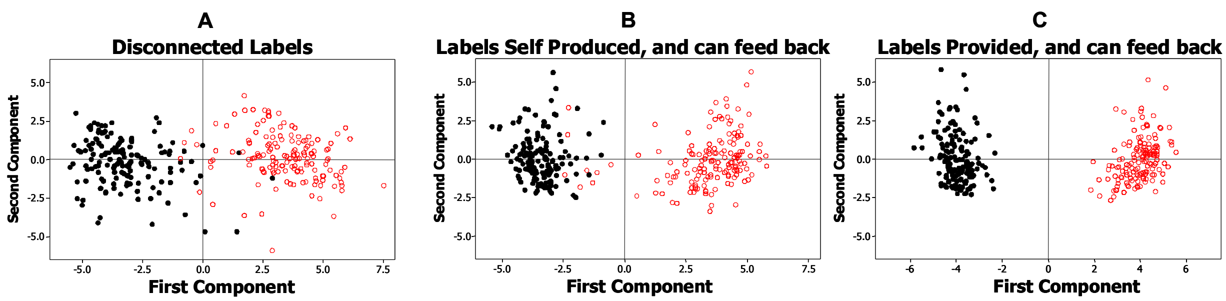

Figure 5: Principal Components Analyses of hidden-unit representations of category exemplars from 2 separate categories. Category structure is enhanced when labels are allowed to feed-back, but at the cost of representing idiosyncratic features of a stimulus.

The simple model presented in Figure 4 can be extended to tasks that require the isolation of specific features. In Section 2.4 I argue that performance on these tasks is causally related to language impairments in aphasia. The feedback of language on conceptual representations occurs not only for whole objects, but also their features and dimensions. Figure 5 shows just the first 2 principal components of the representational space—the actual representational space is high-dimensional with some dimensions or combinations of dimensions likely corresponding to particular dimensions and features of objects. Feedback of labels (e.g., property labels such as "square" "red" or "large") is predicted to have the same kind of cohering effect—facilitating the grouping of objects by their dimensions. This role of labels in realigning representations also suggests a possible explanation why providing labels to young children facilitates performing relational judgments (Gentner & Loewenstein, 2002; Kotovsky & Gentner, 1996; Ratterman & Gentner, 1998).

3.3 Effects of Labels on Visual Processing. Motivations and Mechanisms

3.3.1 Why Vision?

As illustrated by the model above, feedback from category labels can affect not just the "conceptual" representations represented by the hidden-unit layer, but also modulate the lower-level "perceptual" layer that feeds into it.[10] The visual domain offers many possibilities for exploring these types of effects. Perhaps the most important reason for looking for effects of labels on visual processing is that the names (labels) of common objects are associated with perceptual, and in particular *visual* features. The word "chair" may bring to mind certain functional features, to be sure, such as its usefulness for sitting in, but arguably it is the visual features that are most salient. This is not to say that other perceptual modalities are unimportant. Certainly, words like "bark" may activate more auditory than visual representations. Indeed, hearing words like "kick" and "punch" even seems to activate regions of motor cortex corresponding to leg and hand areas, respectively (Buccino et al., 2005). Much remains to be done in exploring the specific influence of verbal labels on representations in different modalities (auditory, motor, somatosensory, olfactory). As a start, the present work looks at the visual domain. A further advantage of examining the visual system is that much is known about its anatomy and physiology and there exist a number of highly robust and simple paradigms for

[10] In the network schematized in Figure 4, conceptual representations are associated with the middle "hidden" layer. This should not be taken as an assumption that there is a dichotomy between perceptual and conceptual representations. Rather, one can view them as a continuum. Neural activity that is primarily driven by bottom-up sensory information corresponds to perceptual representations, while neural activity primarily driven by higher-level (typically more frontal) areas, and occurring even in the absence of sensory input, can be thought to correspond to the conceptual representations. Due to the complex interplay between bottom-up and top-down processing purely perceptual representations may not exist in primates (Mesulam, 1998).

studying effects of various factors on visual processing (e.g., visual quality, cue type, cross-modal interaction, etc.). The section that follows briefly reviews the aspects of the chosen paradigm—visual search—that make it particularly suitable for exploring effects of language on visual processing (see also 4.2). I then review the neurophysiology of the visual system highlighting some mechanisms by which labels may affect visual processing.

3.3.2 The Visual System Hierarchy Revisited

The cornerstone claim of the LFH is that labels modulate lower-level representations through *feedback connections*. To understand the importance of feedback connections in visual processing, it is useful to review the flow of information in the visual system as traditionally conceived:

Visual signals start in the retina, begin cortical processing in V1 and following V2, split into the ventral ”what“ and dorsal ”where“ (or ”how“) streams. The where stream proceeds through MT to MST, LIP, and greater parietal cortex. The ventral ”what“ stream proceeds through V4 to the temporal cortex (areas IT and beyond) (Mesulam, 1998). Mapping studies of connectivity patterns in the visual system of the macaque (Felleman & Van Essen, 1991) report 10 levels of cortical processing. As one follows the anatomical hierarchy from lower to higher cortical regions, receptive fields (RFs) grow in size and complexity (Maunsell & Newsome, 1987): V1 cells are small and are selective for simple features like oriented lines. Cells in IT have large and complex RFs responding to, for instance, hands, or flowers (Desimone, Albright, Gross, & Bruce, 1984).

With few exceptions (Lamme & Roelfsema, 2000; Lamme, Super, Landman, Roelfsema, & Spekreijse, 2000; Spivey, 2007), most theories of visual processing make strong assumptions regarding the consequences of the hierarchical organization of the visual system, for instance, it is common to assume that lower-level areas participate in bottom-up ”sensing,“ while higher-level areas participate in ”perception“ and are influenced by top-down cognitive processes such as goals and motivations (Pylyshyn, 1999). This assumption is made plain in articles with titles such as ”Climbing the cortical ladder from sensation to perception” (Treue, 2003).

On this traditional account, the outputs of lower-level regions comprise the building blocks from which higher-level regions build object representations either in a distributed or in a more localist fashion (e.g., through ”grandmother” (or Jennifer Anniston) cells (Quiroga, Reddy, Kreiman, Koch, & Fried, 2005).

Based in part on studies comparing effects of attentional deployment or “interpretation” of physically ambiguous stimuli, the main assumption has been that as one ascends through the visual hierarchy, there is a shift from neural representations dominated by the sensory input to representations emphasizing the input's perceptual interpretation and the individual's behavioral state (e.g., Treue, 2003). There are two problems with this conclusion. First, the anatomical hierarchy of the visual system is not as strict as assumed by traditional accounts. Second, in systems characterized by high levels of recurrence (through feedback connections), anatomical hierarchies need not translate to functional hierarchies. As I argue next, the visual system is indeed characterized by massive recurrence.

Vision begins with a “feedforward sweep” driven by basic intrinsic features of the stimulus such as contrast, motion, size, and color. Initial cortical processing takes place in primary visual cortex (V1). A strict feedforward model of the visual system would predict steadily increasing response times to visual stimuli as one moves from primary to higher visual regions. Indeed, the earliest response latencies in the monkey evoked by visual stimuli occur at 35 ms in V1, 54 ms in V2, and 61 ms at V4. In direct contrast to this evidentially hierarchical processing, neurons in

MT start responding at 39 ms (MT receives direct inputs from both magnocelullar V1 pathways and subcortical inputs from superior colliculus and pulvinar). MT in turn projects directly to the frontal eye fields (43 ms), and prefrontal cortex (51 ms). We thus see the prefrontal cortex being activated before V2. Temporal regions, with sparser connections to occipital cortex, start responding in the 60-80 ms range, though notably, some (e.g., TPO, TAa) respond several milliseconds before V4 (Lamme & Roelfsema, 2000 for review and meta-analysis).

This feedforward (bottom-up) sweep is complete within 100 ms at which time all the levels of the visual hierarchy are active, from occipital to prefrontal cortex. The inter-level latencies of about 10 ms leave little time for horizontal connections and no time for feedback processes (Tovee, 1994). The RFs of neurons during the feedforward sweep reflect the feedforward connectivity pattern (the so-called classical RFs): small and simple for V1, large and complex for IT (as described above).

Logically, the extremely early activation of areas such as MT, may enable its feedback to affect early visual processing. In fact, because the parvocellular projections to V1 are much slower than magnocellular projections to V1, feedback from MT can reach V1 in time to modulate its activity to the stimulus that caused the feedback in the first place (Hupe et al., 2001; Vidyasagar, 1999). Hupe and colleagues found that interrupting feedback from MT through cooling alters visual responses of V1 neurons within 10 ms of their initial response (Hupe et al., 2001). The authors speculated that the extremely rapid conduction velocities of the V1↔MT pathway mean that some signals from MT can be transmitted (back) to V1-V3 in as little as 1-2ms. Such is the non-correspondence between the anatomical and temporal hierarchy.

What is the role of this feedback activity? Cortical neurons continue to be active after the feedforward sweep and their RFs change dramatically. In V1, cells are retuned from reflecting simple orientation feature within classically small and simple RF, to reflecting figure/ground relationships over a much larger area in as short as 100 ms after onset. As a comparison, V1 cells show orientation selectivity for textures composed of oriented lines (as opposed to single lines) at about 55 ms (Lamme, Rodriguez-Rodriguez, & Spekreijse, 1999).

Modulation of V1 neurons is also seen in more complex tasks. In a task that requires subjects to respond to whether a probe is on a target or distractor curve—performed by mentally tracing the curve from the point of fixation to the probe—RFs of V1 neurons show a response enhancement of 30% when their RF falls on a target curve than when the same curve is a distractor. This modulation appears around 230 ms after onset (Roelfsema, Lamme, & Spekreijse, 1998). Finally, Lee & Nguyen (2001) found that neurons in the primate early visual cortex (V1 and V2) respond to illusory contours. When Kanizsa figures were presented such that their illusory contours were in the receptive field of particular V1 or V2 neurons, those neurons responded. The firing pattern was delayed compared to when an actual (non-illusory) contour was shown, evidence that the neuron was being driven by feedback connections. When the vertices of the Kanizsa figure were rotated so that the illusory figure was no longer "visible," the neurons stopped responding. Responses to illusory contours began about 100 ms and peaked at 125 ms after stimulus onset.

3.3.3 Feedback Activity: Just lights on the screen?

Demonstrations that early visual areas are modulated by context and experience (Chelazzi, Miller, Duncan, & Desimone, 1993; Moran & Desimone, 1985; Pettet & Gilbert, 1992) may seem an epiphenomenological curiosity—the neural version of lights on a computer screen—a signal that is clearly task-relevant, but not-causal of the underlying process. Logically, finding that neurons in V1 are retuned to reflect figure/ground separation does not mean that these later responses participate in either awareness or in producing the behavioral response (Pylyshyn, 1999). Indeed, it is perhaps more parsimonious to assume that because the areas from which the feedback signals arise are already responding selectively to figure and ground, the subsequent activity of the lower areas is redundant (Lamme et al., 2000 for discussion).

Experimental evidence for the causal role of feedback activity requires the ability to manipulate it, separating responses based on feedforward processing from responses based on feedforwards+feedback responses. This is extremely difficult as feedback pathways are anatomically contiguous with feedforward pathways and thus cannot be selectively inactivated or lesioned. The temporal separation between the two processes: feedfoward (<100 ms) and combined feedforward + feedback (>100 ms), also provides little hope of separating out these effects in a purely behavioral task. With the possible exception of very simple tasks that allow the subject to prepare a motor response ahead of time (e.g., probe detection), responses for even the simplest visual decision tasks are sufficiently long for reentrant processes to affect the response. For instance, RTs in the simplest visual search tasks exceed 400 ms. A search for a red element among green elements takes approximately 425 ms (e.g., Wolfe, Butcher, Lee, & Hyle, 2003). Considering that simple RTs to a visual stimulus for college age adults are in the 180-200 ms range (Welford, 1980), the >200 ms that remain constitute an interval during which rich reentrant processing can take place. In addition, even if the first trial of a very fast simple visual task reflect purely bottom-up stimulus characteristics, the numerous trials that characterize visual experiments mean that a bias (a kind of sustained priming) can quickly build up. Williams and colleagues found that when presented with a moving stimulus that could be interpreted as moving in one of two directions (a type of Barber Pole stimulus), neurons in the LIP not only predicted the monkey's eventual response (i.e., its interpretation of physically ambiguous motion), but began to show a motion bias even before the stimulus appeared (Williams, Elfar, Eskandar, Toth, & Assad, 2003). Thus, even random and unpredictable trials appear to be non-independent of previous experience, and thus open to top-down influences.

The conclusion that that activity that emerges in lower visual areas following horizontal and feedback influences is causally involved in visual awareness is supported by at least two types of evidence. First, application of transcranial magnetic stimulation (TMS) to early visual cortex disrupts visual awareness not only during the feedforward stage (-30 ms to +50 ms relative to stimulus onset), but also between 80 ms and 120 ms—time at which processing has already reached the highest cortical areas (Corthout, Uttl, Walsh, Hallett, & Cowey, 1999). Thus it appears that interrupting activations in early visual areas produced by feedback activity interferes with visual awareness. In a further demonstration of the causal relationship between V1 activity and awareness, applying TMS to V1 has been found to inhibit participants' ability to perform mental imagery tasks, which on a traditional hierarchy account should not require the participation of lower-level visual areas at all (Kosslyn et al., 1999; Kosslyn, Ganis, & Thompson, 2001).

A second source of evidence that implicates reentrant processes in visual awareness is the curious phenomenon of common onset masking (Di Lollo, Enns, & Rensink, 2000). Masking by common onset consists of the stimulus and a mask appearing simultaneously. This is possible by using a metacontrast mask (a contour of a greater contrast that closely surrounds, but does not occlude the object). The object and the mask are presented briefly together. As it turns out, if the mask remains visible after the object inside it has disappeared, the object is wiped from awareness (as long as there is initially more than one object on the screen). If the stimulus and mask disappear together, participants have no trouble reporting the identity of the object.

Common onset masking works even when an object-substitution mask is used comprising of several dots that surround the object. As long as the mask (which is now just a few dots) remains on the screen after the object has disappeared, participants' ability to report the identity of the object is dramatically decreased. As the number of objects that appear simultaneously increases, performance for reporting the identity of the masked object drops to chance. At the same time, their ability to report the object if it disappears together with the mask is unaffected. Di Lollo and colleagues (2000) explain this phenomenon by assuming that awareness requires a match between top-down and bottom-up activity. When an object is processed with a mask, they are essentially processed as a single unit. If the object then disappears, the feedback activity (object+mask) making its way back to early visual cortex now conflicts with the bottom-up activity still present (mask only). This bottom-up activity overwrites the pattern that is feeding back. Though this account may seem odd, it is consistent with TMS studies that find that disruptions of activity in early visual cortex >100ms after stimulus onset disrupt visual awareness. Common onset object substitution masking has resisted attempts at explanations that do not involve recurrent connections although Macknik and Martinez-Conde have recently argued for a mechanism relying on feedforward and horizontal projections (i.e., lateral inhibition circuits) (Macknik & Martinez-Conde, 2007).

Why should activity in lower visual areas be causally related to visual awareness? This goes beyond the scope of this thesis. One possibility is that bottom-up input is initially too chaotic, reflecting both relevant and irrelevant stimulus properties. This activity is cleaned up by context- and task-sensitive feedback input that joins together disparate features of a stimulus, and nudges the visual system into a state of coherence (Serences & Yantis, 2006).

3.4 *Visual Processing and the LFH*

In the last section I presented evidence implicating feedback connections in processes from figure/ground separation to mental imagery. The LFH proposes that just as perception of illusory contours (Lee & Nguyen, 2001) or objects like the Necker Cube (Rumelhart et al., 1986) is best described as an interactive, parallel process, so perception of familiar objects is hypothesized to be interactively influenced by verbal labels with which they are associated. Feedback connections between words (lexical units, or verbal representations more broadly) and lower-level semantic/perceptual representations are in fact present in a number of word-production models (Dell, 1986; Dell & Oseaghdha, 1992; Dell, Schwartz, Martin, Saffran, & Gagnon, 1997; Plaut, 2002). The role that these projections play in forming and maintaining semantic or perceptual representations has never, to my knowledge, been investigated.

Because labels are differentially associated with object features (depending on the degree to which the feature is diagnostic and typical to the named category), the activation of a label is predicted to have the overall effect of making objects more typical, facilitating their processing. This may be achieved by cleaning up perceptual noise and/or through highlighting of diagnostic dimensions/features while collapsing representations of irrelevant differences (e.g., see Experiment 3a). In either event, hearing an object label is predicted to facilitate the visual processing of exemplars from the named category, possibly in parallel and throughout the visual scene. The final section tests this prediction using the paradigm of visual search.

4 Using Visual Search to Explore Linguistic Influences on Perceptual Processing

4.1 Visual Search: A Short Primer

The visual search paradigm involves finding a target among distractors (also called non-targets). The target is either known ahead of time, or not, in which case it can be located by its unique status (i.e., if there is a unique item in the display, that is the target). Most studies involve finding a single target rather than one of several targets, thus there is at most one target per trial. Typically, 50% of trials contain the target (target-present or positive trials) and the rest contain all non-targets (target-absent or negative trials). Participants are instructed to respond *present* or *absent* as quickly and accurately as possible. The critical dependent variable is the search slope—the amount of increase in RT for each additional display element, measured in ms/item. Errors tend to correlate with difficulty of search and so tend to rise with longer RTs. Too sharp a rise in errors with increasing difficulty, however, signals a speed-accuracy tradeoff.

The visual search paradigm has become instrumental in studying the mechanisms of attention and visual processing following the publication of the feature integration theory of attention (FIT) (Treisman & Gelade, 1980). Treisman and Gelade drew a qualitative distinction between search for "basic features" such as color, basic shape, and orientation, which can be achieved preattentively, and thus in parallel. The diagnostic for a parallel search is a search slope of less than 10 ms/item (often 5 ms/item or below). Preattentive processing of basic features means that one can detect such a basic feature without having to serially deploy attention to each item in the display. This shows up in parallel search slopes and is often correlated with the phenomenological experience of target "pop-out," occurring, for example, when one searches for a red thing among green things (Treisman & Gormican, 1988). In contrast, when search involves a conjunction of features, attention needs to be deployed serially. For instance, when searching for a red T among blue T's and red X's, the two features (shape and color) have to be integrated before a target can be located. According to FIT, this integration requires attention, which is serially applied, requiring the examination of all items until the target is found or all items have been examined (Townsend, 1990 for review).

In contrast to the serial/parallel dichotomy of FIT, a number of theories have proposed that all search has a parallel component. What determines the efficiency of a search task is not whether a target can be found preattentively, but rather the perceptual difference between the target and non-targets as well as the perceptual homogeneity of non-targets. For instance, in Duncan and Humphrey's competition theory (1989) the target and non-targets compete for attention. The competition can be resolved more quickly when the perceptual difference between them is greater and when the non-targets are more perceptually homogeneous allowing them to be grouped (and inhibited) by the visual system. Guided Search Theory (Wolfe, 1994) similarly rejects a serial/parallel dichotomy, arguing that search proceeds through a top-down guidance of feature maps that are derived from bottom-up stimulus information. For instance, searching for a red vertical line among red horizontal lines and green horizontal lines involves a conjunction, and thus according to FIT should be serial. Yet, this search can be efficient as long as target-identity is known beforehand (e.g., Wolfe & Horowitz, 2004; Wolfe, Horowitz, Kenner, Hyle, & Vasan, 2004). What happens according to Guided Search is that the scene is filtered in parallel through feature channels (one channel for color, another for orientation). The outputs of these channels produce feature maps (e.g., a color map and an orientation map). Top-down commands then activate particular values, for instance, "red" in the color map and "vertical" in the orientation map. This is the guided search part of the theory. The top-down guidance results in a map of locations of red things and a map with locations of vertical things. Combining these maps produces a master activation map. The purpose of this master map is to direct attention. In the absence of any contrary endogenous cues, attention is deployed to the location with the peak activity, followed by the next highest peak, and so on. In an efficient search (i.e., a search with a very shallow slope), the target item produces the highest level of activation regardless of the number of non-targets. It is therefore the first item to attract attention. Efficient *guided* search is possible as long as the target is defined by basic features (even if they have to be combined). Although neither the color or orientation maps contain the information necessary to find the target, the peak of the combined map corresponds to the red-vertical target. In this framework, *inefficient* search results when the target is defined by non-basic features which by definition cannot guide the deployment of attention. Some examples of basic features are color, motion, orientation, size, luminance, and curvature. Some examples of features that arguably cannot guide attention are optic flow, color-change, semantic categories, letter identity, novelty, faces, and intersections (Wolfe & Horowitz, 2004 for review). So, although all search has a parallel component, some features and feature combinations are able to affect the deployment of attention more effectively than others, leading to shallower slopes (Wolfe, 2003; Wolfe & Horowitz, 2004).

One of the most prominent features of performance on visual search tasks is that inefficient search produces extremely linear search functions (i.e., RTs increase proportionally to the number of items). If all search is essentially parallel, why do “inefficient” searches show a linear increase in time with linear increases in the display size? The most striking demonstration of how linearly increasing search times emerge from a parallel process is Spivey’s normalized recurrence model (Reali, Spivey, Tyler, & Terranova, 2006; Spivey & Dale, 2005). In this model objects compete for "targethood." Object features are integrated in parallel until a criterion is reached. With increasing set sizes, multiple units sum their activation (which is the normalized) to achieve a criterion of activation at approximately linear intervals in time. Within this framework, parallel search is produced when a minimum number of competition cycles are required for an object to cross the target threshold. As the number of elements increases, so does the number of iterations (insofar as search requires integrating multiple features). Remarkably, the required number of iterations increases linearly with the display size. Such a model provides an existence proof that linear functions can be produced by a system in which multiple partially active representations are simultaneously competing (Spivey & Dale, 2005, p. 121).

Despite making much progress in understanding both the deployment of visual attention and the mechanisms governing performance in visual search tasks, Wolfe's 1998 dictum that no current model can account for all the major visual search findings (Wolfe, 1998) remains true today.

4.2 Categories and Visual Search

In competing for attention during visual search, not all features are equal. Targets defined by some features such as color, motion, orientation, and size can guide the deployment of attention. Targets defined by features such as intersections, semantic category, and novelty arguably cannot (Wolfe & Horowitz, 2004 for review). This deployment of attention corresponds to a warping of the attention saliency map in such a way as to make the target have a sufficiently higher saliency than of the distractors (non-targets). In addition to the most basic features like color, orientation, and size, attention can be guided by features with some complexity such as curvature, pictorial depth, number, aspect ration glossiness, and topological status: search for targets defined by these features is efficient and can survive some degree of distractor heterogeneity (Wolfe & Horowitz, 2004 for review).

One long-standing debate has been whether it is possible to guide attention to particular object categories. That is, can attention be guided by conceptual (or semantic) status of targets. Alternatively, does the visual processing involved in search merely reflect the physical properties of the stimuli such as the physical similarity between targets and non-targets (Theeuwes, 1993)?

Objects in preattentive vision have been argued to be "shapeless bundles" of features (Wolfe & Bennett, 1997) that are loosely organized and transitory (Rensink, 2000). On this account, it may be surprising to find that that search among complex objects can in fact be very efficient. For instance, Levin et al. (Levin, Takarae, Miner, & Keil, 2001) reported that searching for an artifact among animals is achieved with a search slope on the order of 5.5ms/item. A conclusion that this efficiency is produced by the categorical status of the target and non-targets turns out to be unwarranted. In the same paper, Levin et al. (2001) showed that efficient search for an artifact among animals can be distilled to perceptual rather than conceptual differences—pictures of man-made artifacts tend to be more rectilinear than pictures of animals—and it is this difference that is probably responsible for the efficient search. Similar perceptual-based explanations have been invoked to explain the classic category effects such as the alphanumeric effect, the finding that cross-category searches (searching for a number among letters or a letter among numbers) are faster than within-category searches (searching for a number among numbers or a letter among letters) (White, 1977; Karlin & Bower, 1976; Duncan, 1983). With single predefined targets, the category effects vanish when perceptual differences among numbers and letters are controlled (Krueger, 1984) [11]. Why this occurs is still unknown. The well-known finding that

[11] Category effects do re-emerge, however, when participants are asked to search for one of many targets. So, when asked to report whether one of several letter or number targets exists, participants are faster to search among

search is faster for an O ("oh") among numbers and a perceptually identical target O (called "zero") among letters—the so-called oh-zero effect (Jonides & Gleitman, 1972)—has failed multiple attempts at replication (Duncan, 1983; White, 1977) (though see Taylor & Hamm, 1997 for an oh-zero effect in an attentional blink paradigm).

An important clarification is the difference between the current conception of category and the loose definitions applied in much previous work. Following Rosch (1973), a category comprises a collection of objects that share certain features, and often has a prototypical structure. A label is most naturally applied at the basic level, so while "poodle", "dog", and "animal" are valid category descriptions of the creature, "dog" is the one applied most easily. The experiments in this paper are most concerned with subordinate and basic level categories because it is here that categories are most likely to have perceptual features in common. While a perceptual prototype of a chair will look like a chair, a perceptual prototype of *furniture* is probably less meaningful, and indeed such superordinate labels do not offer much guidance in visual search (Wolfe et al., 2004). In most early studies, categorization was defined at the superordinate level—classifying a stimulus as a letter or number, while identification was defined at the basic level—"identifying" a stimulus as a 'b'(Duncan, 1983; White, 1977). The failure to find category effects in visual search controlling for perceptual confounds therefore may have been due to the use of superordinate categories, the members of which often have little in common and so would have little perceptual coherence. This confusion between levels persists in contemporary studies. For instance, Levin et al. (2001) found that jumbling animal and artifact stimuli slows search for an artifact or animal only minimally. This was taken as evidence against the possibility that "search is based on a completed structural code" (Levin et al. 2001 p. 683).

---

distractors from a different category than from the same category even controlling for perceptual confounds. Why this occurs is still unknown.

But while jumbling up the features of an elephant makes it more difficult to identify it at the basic level (an elephant), it is still trivial to categorize it at the superordinate level, easily distinguishing between animals and artifacts).

While effects of conceptual categories in visual search have had a checkered history, there is clear evidence of effects of categories that have been argued to be perceptual nature (but see below). For instance, Wolfe and colleagues (Wolfe, Stewart, Friedmanhill, & Oconnell, 1992) found that when searching for oriented lines among differently-oriented distractors, search proceeds as though there are perceptual categories of vertical, horizontal, steep, shallow, right, and left—search within these categories is slower than between even though the perceptual distances remain constant. Daoutis et al.(2006) found these effect with color categories with controlled perceptual distances. The Guided Search model (Wolfe, 1994) elegantly accounts for these findings through top-down effects of "channels" that have peak activations at the centers of these categories (i.e., the vertical channel responds maximally to vertical lines). What remains unknown is where these channels come from and how they come to encode such information as steepness and left-ness. While it is possible that such fundamental perceptual properties as verticality arise from ecological distributions of lines and the early receptive fields, other hypothesized channels seem more arbitrary (e.g., steep, shallow, and purple channels). Might categorical effects be produced in part by linguistically derived categories? That is, might the ability to use the label "steep" facilitate selection of the steep elements, perhaps by enabling more efficient grouping among elements that comprise the "steep" and "shallow" categories. Some evidence comes from a study showing that learning arbitrary associations between tilted lines and arbitrary labels produces such categorical effects. For instance, learning to associate

+45° and 90° lines with the label "pencil" and -45° and 0° lines with the label "elephant." results in faster search when targets and non-targets span the category boundary than when they are both in the same category (Smilek, Dixon, & Merikle, 2006). This finding, however, is subject to some of the same familiarity confounds discussed in Section 1.4.

A additional potential source of evidence for conceptual effects in visual search comes from search asymmetries. I say potential because to date these effects have been framed in terms of differences in perceptual coding as brought on by familiarity, rather than any effects of concepts or meaning. In Treisman's original formulation (Treisman & Gormican, 1988), it was hypothesized that it is easier to detect a "deviant" among standard stimuli than to find the standard stimulus hiding among deviants because a deviant is defined by a presence of a feature, while a standard stimulus is defined by an absence, making it harder to locate. For instance, finding Q (defined by the presence of a bar) among O's, is easy:

OOOOOOQOOOOO

The reverse, is more difficult:

QQQQQQQQOQQQ

Similar asymmetries, more difficult to explain, are found when searching for a unfamiliar (И) among familiar N's:

NNNИNNN

which is very easy (0-6 ms/item). The reverse, finding an N among И's is much more difficult:

ИИИИИИN

and produces slopes 60-80ms/item. This search asymmetry, albeit less extreme, is found with a variety of letter and even pictorial stimuli. For instance, it is faster to find a outline of an upside down elephant among upright elephants (5 ms/item) than vice versa (17 ms/item) (Wolfe, 2001).

This simple effect has baffled researchers. Following FIT, work on "novel pop-out" (Johnston, Hawley, & Farnham, 1993) suggested that novelty itself is a basic feature and its presence guides the deployment of attention to the novel object like the mirror-reversed N or the upside-down elephant—a view supported by Wang, Cavanagh, and Green's (1994) study showing that the presence of a novel letter among standard letters leads to efficient search, but a search for a standard letter among other standard letters does not. The authors concluded that "familiarity itself might be considered a primitive feature which can be processed preattentively." (c.f., Wolfe & Horowitz, 2004). Malinowski and Hubner (2001), on the other hand, argued that rather than novelty being a feature, the letters themselves may have become basic features, so that a search that involves searching through familiar letters (which are by their account, basic features) is easier than a search through unfamiliar features. In their study they showed that for subjects familiar with both the regular N and its mirror image—И—due to knowledge of both the Roman and the Cyrillic alphabets, the asymmetry disappears. A more detailed discussion about familiarity effects can be found in 4.7.4. The most parsimonious explanation to explain familiarity-based search asymmetries has been offered by Shen and Reingold (2007) who showed that what matters most is distractor familiarity. Search through unfamiliar distractors is harder than search through familiar distractors. Perhaps because there is at most one target, but multiple distractors, target familiarity does not significantly affect search performance. In all these studies perceptual familiarity is confounded with meaningfulness. Thus, we know that meaningful/familiar items seem to be easier to process than meaningless/novel items. Explanations have focused on perceptual novelty. So, someone not familiar with the Cyrillic alphabet has not seen the symbol И very often. An alternative (and to my knowledge, untested) view is that in addition to differences in novelty, processing of unfamiliar items is more difficult in part because they are meaningless, that is, their processing is not aided by feedback from top-down category-level representations.[12] The idea that ascribing meaning to otherwise perceptually novel stimuli may facilitate their processing hypothesis is discussed in further detail and tested directly in Experiment 2 (sections 4.7).

If category-membership rather than perceptual familiarity mediates performance (with more meaningful / easier to categorize items processed more quickly), then according to the LFH, categorizing items further by referring to them with the linguistic category should produce more efficient search than presenting the items alone. While participants may automatically categorize familiar and meaningful items, referring to them by name should facilitate their processing through the same mechanisms that result in faster processing of N's than И's. Testing this hypothesis is a primary goal of the studies presented in the subsequent sections.

4.3 Do Labels Deploy Categories?

Clearly, attention can be guided by the informational content of language—being told to search for a red-vertical initiates search for a target with these and not other attributes (Wolfe, 1994; Wolfe et al., 2004). The first evidence that on-line language processing alters behavior in visual search beyond this simple informational aspect came from Spivey (Spivey et al., 2001) who showed that search defined by a color/orientation conjunction is made efficient by language

[12] Some indirect support for this position comes from a study by Levin and Angelone (2001) in which the authors found that distance from category prototype played a significant role in producing a search asymmetry with typical distractors being rejected more quickly than atypical distractors. Typicality may be predicted to impact visual processing if the latter is mediated by category membership (meaningfulness). Alternatively, typical stimuli are just more perceptually familiar than atypical stimuli.

presented concurrently with the search display. When spoken instructions (e.g., "find the red vertical") were timed such that the target information ("red vertical") was delivered concurrently with the search display, search slopes were nearly horizontal. When the identical utterance was presented prior to the appearance of the search display, the slopes were on the order of 15 ms—typical for this type of conjunction search. To explain the effect, the authors argued that the serial delivery of the feature information allowed participants to divide a conjunction search into two simpler feature searches—objects would first compete on the color dimension, and then on orientation.[13] Reali et al. (2006) extended this finding from a blocked to a mixed design and modeled the results using a normalized recurrence connectionist network mentioned in Section 4.1.

Also relevant are studies showing that labels can produce *shifts* of spatial attention. Spatial words such as "above," "below," and "next to" have been shown to shift attention in the corresponding direction as measured by faster RTs to a probe appearing in those locations (Carlson & Logan, 2001; Logan, 1995). The labels bias attention even when the terms are entirely non-predictive of target location. Interestingly, this attentional bias is as strong as when using visual cues like arrows (Hommel, Pratt, Colzato, & Godijn, 2001). Although the relative salience of spatial terms compared to visual cues has been debated (Gibson & Kingstone, 2006), it is clear that familiar spatial terms can swiftly guide visual attention.

What is not known is whether linguistic effects on visual search are more than just a matter of additional information and information delivery. Recall that identifying an object as a *something* involves categorizing it. Because labels refer to categories, labeling an object may

[13] Because overall RTs were much longer in the concurrent condition, the data are also consistent with a possibility of a ceiling effect on these trials with the shallow slope being an artifact of longer RTs in small displays, though unreported pilot studies possibly rule out this interpretation (Reali, pers. comm).

make its representation more categorical, in effect making it a "better" (or at least a more typical) object. The LFH predicts that a label, through top-down feedback, activates visual features typical to the named category. One consequence may be that labeled objects become more "coherent," resulting in more efficient search. This effect should occur even when the label is not an informative cue, that is, when target identity is already known. To the extent that the effects of the label are automatic and not strategy dependent, they be observed in a within-subject paradigm. To the extent that they are temporally transient, as would be predicted from a feedback-driven effect, its effects should be observed in a paradigm that intermixes *label* and *no-label* trials. That is, facilitation induced by the label on trial *n* may not transfer to the no-label trial *n*+1. Furthermore, insofar as the presentation of the label produces effects throughout the visual scene, not only should naming the target facilitate performance, but naming the non-targets (e.g., ignore X) may also produce facilitation, for instance by producing stronger perceptual grouping of the now more typical non-targets. In other words, while overt naming is a serial process in that we label one object at a time, the *effects* of the label on visual processing may occur in parallel, affecting numerous objects throughout a visual scene. A further question (addressed in Experiment 2) is what happens when labels *do* add information to the display, as when a perceptually novel figure is made meaningful by referring to it with a familiar category label.

4.4 Research Questions

Because labels denote categories, hearing a label may modulate perceptual processing such that the perceived objects are more influenced by the named category—the Language Feedback Hypothesis. The aim of the experiments described in this section is to test this hypothesis by answering the following questions:

1. Does on-line presentation of category labels facilitate visual processing as measured by a visual search task? Does the effect of labels depend on visual typicality of the named items? In Experiment 1a the naming of targets was manipulated while in Experiment 1b the naming of non-targets is manipulated.
2. What is the role of meaningfulness in processing perceptually novel stimuli? In Experiment 2 I investigate a curious finding from Experiment 1a-1b and test the hypothesis that difficulties in processing perceptually novel items may arise from their being meaningless rather than perceptually novel.
3. Are the effects of labels on visual processing mediated by the visual quality of the named stimuli? Experiments 3a-3b test the hypothesis that labels "clean up" visual stimuli, and so their effects should be strongest for objects of low visual quality.
4. What is the relationship between the ability to categorize a stimulus and the effects that hearing the category name has on visual processing? Insofar as labels affect visual processing by inducing more categorical processing of display elements, there should be an intimate relationship between categorization RTs, visual search RTs, and effects of labels. Experiment 3b investigates these links and includes a study of individual differences.
5. Can strengthening the association between labels and visual exemplars facilitate visual processing? Experiments 4a and 4b examine the consequences of a brief category training session on subsequent visual processing.
6. What are the individual differences in effects of labels on visual processing? Experiment 4b includes a correlational analysis of individual factors that contribute to effects of labels on visual processing.

Additional experiments presented in the appendix: *Conceptual Grouping Effects: Categories Matter* (*and named categories matter more)* tackle these additional issues concerning the interactions of categories, labels, and visual processing:

1. Do conceptual categories produce perceptual grouping?
2. Does this effect arise from long-term or on-line influences of categories?

4.5 Experiment 1a: Naming the Target: On-line effects of labels on visual processing

4.5.1 Participants

Twenty-four subjects, 18-22 years old volunteered for the experiment in exchange for course credit. They were naïve to the experimental hypothesis.

4.5.2 Stimuli and Procedure

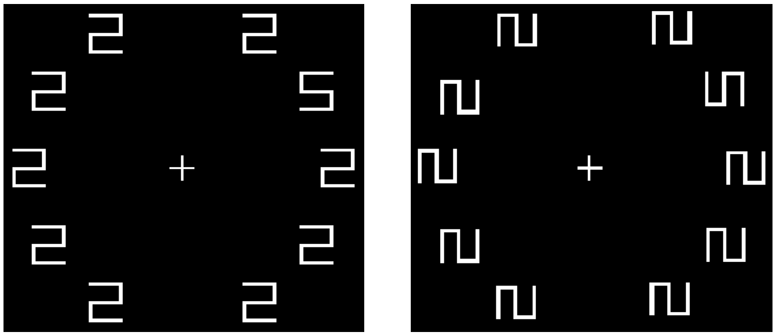

Figure 6: Two sample trials in experiments 1a and 1b. Upright search for a 5 among 2's (left). Rotated search for a 5 among 2's (right).

Four distinct stimuli were used in the search trials: upright numbers: and , and rotated numbers: and . The rotated numbers were included on the assumption that they were less typical instances of the 2 and 5 categories, allowing me to compare the effect of labels on more versus less typical category members. Effects of typicality are explored more thoroughly in Experiments 3 and 4. Participants were told to think of all symbols as the numerals 2 and 5. I used numeric characters in these studies because they are perceptually simple and, being overlearned, have strong category representation and they have also been used in several other studies (e.g., Wang et al., 1994). As is true of most categories, these stimuli can be classified at multiple levels of abstraction—a 2 can be a "number" an "even number" or a "two" (Posner & Mitchell, 1967). Classifying a 2 as "two" can be thought as a basic-level task in the sense that it is generally faster to identify alphanumeric characters at this level compared to more superordinate levels (Dick, 1971; Posner, 1970), probably because a "2" is more frequently classified as a "two" than as a "number."

To assess the impact of auditory labels, a recording of the words "find the" (Experiment 1a) and "ignore the" (Experiment 1b) was spliced with the words "two", "five", and a segment of white noise, creating 3 audio clips: "find the five", "find the two", "find the [noise]" where [noise] corresponds to a white noise of equal intensity and duration to the "five" and "two" clips. Thus, all trials included some auditory stimulation, but only the *label* trials included the category label, here, “two” or “five.” Experiment 4a, 4b, and the Experiment 3 in the Appendix substitute the white noise used in the no-label trials with alternative fillers such as “find the target.”

All auditory stimuli were adjusted to be of the same intensity and length (1000 ms).

The characters were white on a black background and had a visual-angle size of $.7^{o} \times .8^{o}$. The characters were arranged along the circumference of an imaginary circle having a diameter of 7^{o} around a fixation cross ($.5^{o}$ diameter). The placement of the target and distractors was random with the stipulation that the same number of items were present on the left and right sides of the display. Sample trials containing upright and rotated stimuli are shown in Figure 6.

Participants completed two parts in counterbalanced order. In one part they searched for a 2 among 5s. In the other, for a 5 among 2s. Thus, the target and non-targets were always known and remained the same from trial to trial, switching halfway through the experiment. Each part consisted of 24 conditions: target present/absent × display size (4, 6, or 10) × orientation (upright or rotated) × *label* or *no-label*. All the conditions were seen by all the participants. Participants completed 10 blocks for a total of 240 trials searching for 2s and 240 searching for 5s. Each block began with 15 practice trials. The inter-trial interval was 750 ms. Trial order was random with the target present on exactly half the trials. Participants gave 2-alternative target *present* / *absent* responses using a gamepad controller. Participants were instructed to respond as quickly as possible without compromising accuracy. If accuracy dipped below 92% for 24 trials, participants saw a display asking them to try to be more accurate. Feedback in the form of a buzzing sound was provided for incorrect responses. Response mapping (left hand *present* vs. right-hand *present*) was counterbalanced between participants.

Prior to the appearance of each search display, participants heard a sound clip label the targets on half the trials: “find the two” / “find the five”—the *label* condition, or a sound clip in which the target label was replaced by white noise ("find the [noise]")—the *no-label* condition. The search display appeared 600 ms after the end of the sound-clip.

4.5.3 Results

Search performance was analyzed using a within-subject ANOVA with display size, orientation (upright or rotated), and labeling (with-labels, without-labels) as within-subject factors. Errors were higher for the rotated trials (9%) compared to upright trials (5%), $F(1, 23) = 8.30$, $p < .01$. There were no other accuracy effects.

Reaction times for all conditions are graphed in Figure 7. Error bars correspond to within-subject 95% Confidence Intervals (Loftus & Masson, 1994). Reaction time analyses included correct responses only and excluded RTs less than 150 ms and greater than 3 standard deviations above the mean. Analyses will focus on the target-present trials. As can be seen in Figure 7, target-absent data paralleled target-present data, but effects of labels in this and all other experiments are limited to trials in which a visual discrimination between the target and non-targets is required. Consistent with the findings of Wang et al. (1994), RTs were longer on trials that involved searching for the rotated targets, $F(1, 23) = 64.43$, $p < .0005$. Unlike Wang et al's (1994) findings, however, there was no reliable difference between the target-present search slopes for rotated versus upright trials, $(2, 23) = 2.14$, $p > .12$. Search slopes for the target-absent trials were, however, greater for the rotated stimuli than the upright stimuli, $F(2, 23) = 7.42$, $p < .01$.

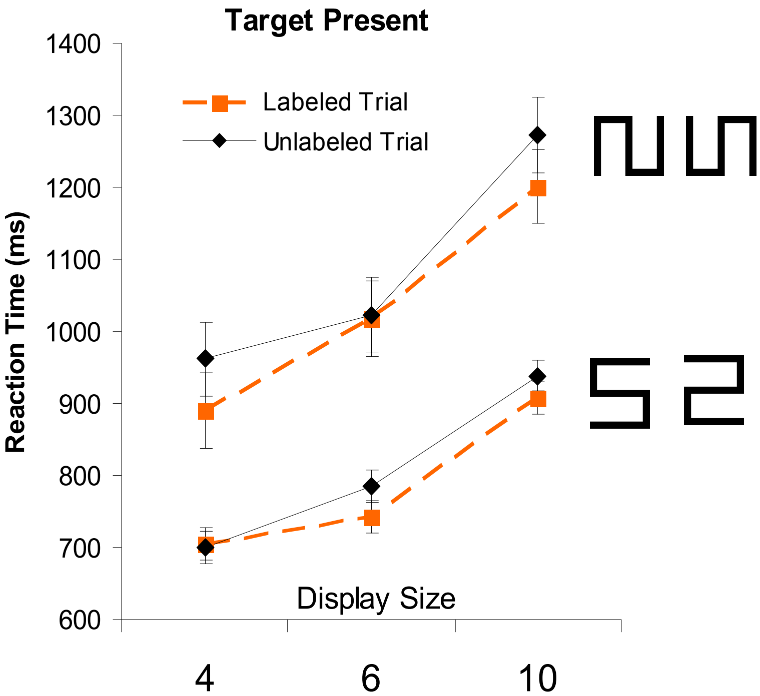

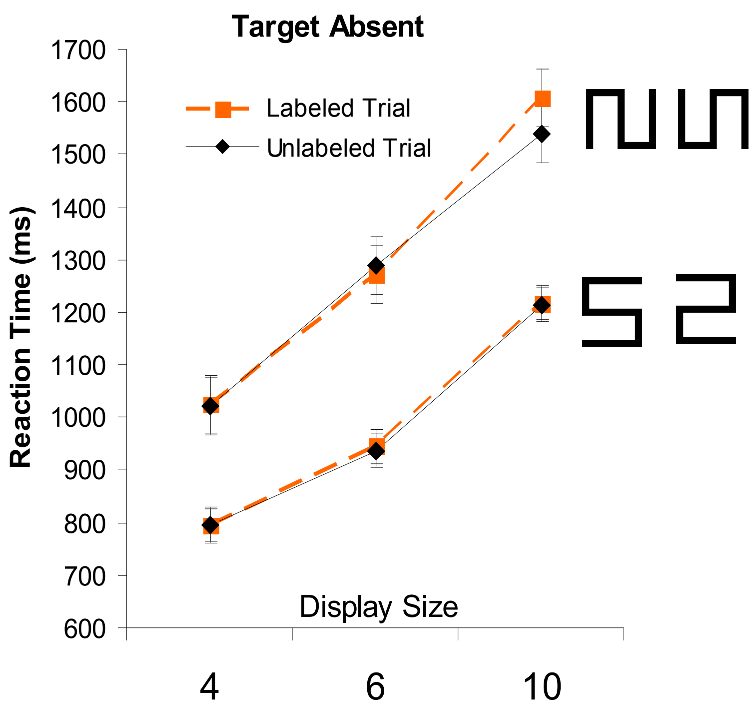

Figure 7: Experiment 1a search times for labeled vs. unlabeled upright and rotated trials. Labels name the target. Labeling and orientation are intermixed (manipulated within each block). Error bars are 95% Within-Subject Cis.

Search on *label* trials was significantly faster than search on *no-label* trials, $F(1, 24) = 5.15, p < .05$. This facilitation reached significance for the upright trials, $F(1, 23) = 5.59, p < .05$, but not the rotated-trials, $F(1, 23) = 2.49, p > .13$. However, the orientation × labeling interaction was not significant.

4.5.4 Discussion

Considering that participants knew the target and non-target identity on each trial (simply because they were held constant from trial to trial), hearing the labels provided no unique information to the participants. Yet, search on trials on which the known target was labeled, was significantly faster than on trials on which it was known, but not labeled. The labeling effect displayed some selectivity, being slightly stronger for the upright compared to rotated trials, suggesting a modulation familiarity / typicality modulation. This finding is explored further in subsequent studies.

As discussed above, search through unfamiliar non-targets is generally less efficient than search through familiar non-targets (Frith, 1974; Malinowski & Hubner, 2001; Shen & Reingold, 2007; 1994). The results of the present study departs from this pattern, revealing a reliable difference in RTs, but little difference in efficiency between the familiar upright, and the unfamiliar rotated stimuli (see also Experiment 1b). One possibility is that labels, although having little on-line effect on the search times of the rotated stimuli nevertheless resulted in a type of sustained priming of the rotated stimuli. Alternatively instructing participants to think of all the search elements as members of meaningful categories (2s and 5s) resulted in facilitated processing for the otherwise unfamiliar symbols ⊓⊔ and ⊔⊓. This latter possibility is tested in Experiment 2.

4.6 Experiment 1b: Naming the Non-Targets: Scene-Wide Effects of Labels

The finding of faster search on trials on which the target was labeled (Experiment 1a) can be explained in several ways. One possibility is that the name made features of the target more salient, facilitating deployment of attention to stimuli matching the named category. A related, but distinct possibility, consistent with the Label Feedback Hypothesis is that the label alters the competition dynamics between the target and the non-targets (Duncan & Humphreys, 1989; Spivey & Dale, 2004). The target is selected more readily when labeled because its visual representation becomes more distinct from that of the non-targets. These two alternatives can be teased apart (albeit partially) by labeling the non-targets rather than the target.

On the first account, labeling the *non-targets* rather than the targets would direct attention towards the non-targets making search more difficult on the labeled than non-labeled trials. On the second account, labeling the non-targets would produce an even greater facilitation because on any search trial there is at most one target, but multiple non-targets whose processing can be affected by the label. Finding a robust facilitatory effect of labeling non-targets would also suggest that labels can exert a scene-wide effect. While it is conceivable that target labels in Experiment 1a constituted a kind of reminder ("two, right, I'm supposed to find the two"), leading participants to more actively search for the item, finding that labeling the *non-targets* can also facilitate performance would rule out this type of explanations because a focus on non-targets following the presentation of the label would impair rather than facilitate performance.

4.6.1 Participants

Twenty-four new subjects, 18-22 years old volunteered for the experiment in exchange for course credit. All were naïve to the experimental hypothesis.

4.6.2 Stimuli and Procedure

The stimuli and procedure were identical to Experiment 1a except the label "find the" was replaced by "ignore" thus labeling the non-targets rather than the targets. There were three sound clips: "ignore five", "ignore two", and "ignore [noise]." Although the labels referred to multiple 5s and 2s, respectively, the words "fives" and "twos" seemed unnatural due to the rare occurrence of plural markers on number words. Hence the singular labels "two" and "five" were used. All auditory stimuli were adjusted to be of the same intensity and length (1000 ms).

As in Experiment 1a, target and distractor identities were blocked. Consequently, participants always knew ahead of time what the target and distractors were going to be—the linguistic label did not tell them anything they did not already know. Contrasting the label and no-label conditions thus corresponded to the contrast between knowing target and non-target identity and hearing the non-targets actually labeled, and having the same prior information about the targets and non-targets, but without hearing the label. As in Experiment 1, participants searched for a 2 among 5s and then for a 5 among 2s, with the order of the two parts counterbalanced. Unlike target and distractor identities which were blocked, orientation and labeling conditions were intermixed within each block.

4.6.3 Results

Analysis of errors revealed a significant effect of orientation, $F(1,23) = 13.82$, $p < .001$, with rotated numbers producing more errors (8%) compared to upright numbers (6%). Labeling did not significantly affect accuracy, $F(1,23) = 2.31$, $p > .13$.

In a further departure from Wang et al.'s (1994) finding of an almost two-fold difference in search slopes between familiar upright numbers and perceptually rotated "numbers," the present experiment yielded longer RTs on trials with rotated numerals, $F(1, 23) = 42.12$, $p < .0005$, but the display-size × orientation interaction was far from significant, $F(2, 23) < 1$. As shown in Table 1, the difference in slopes was only about 1 ms/item. Reaction time results are graphed in Figure 8.

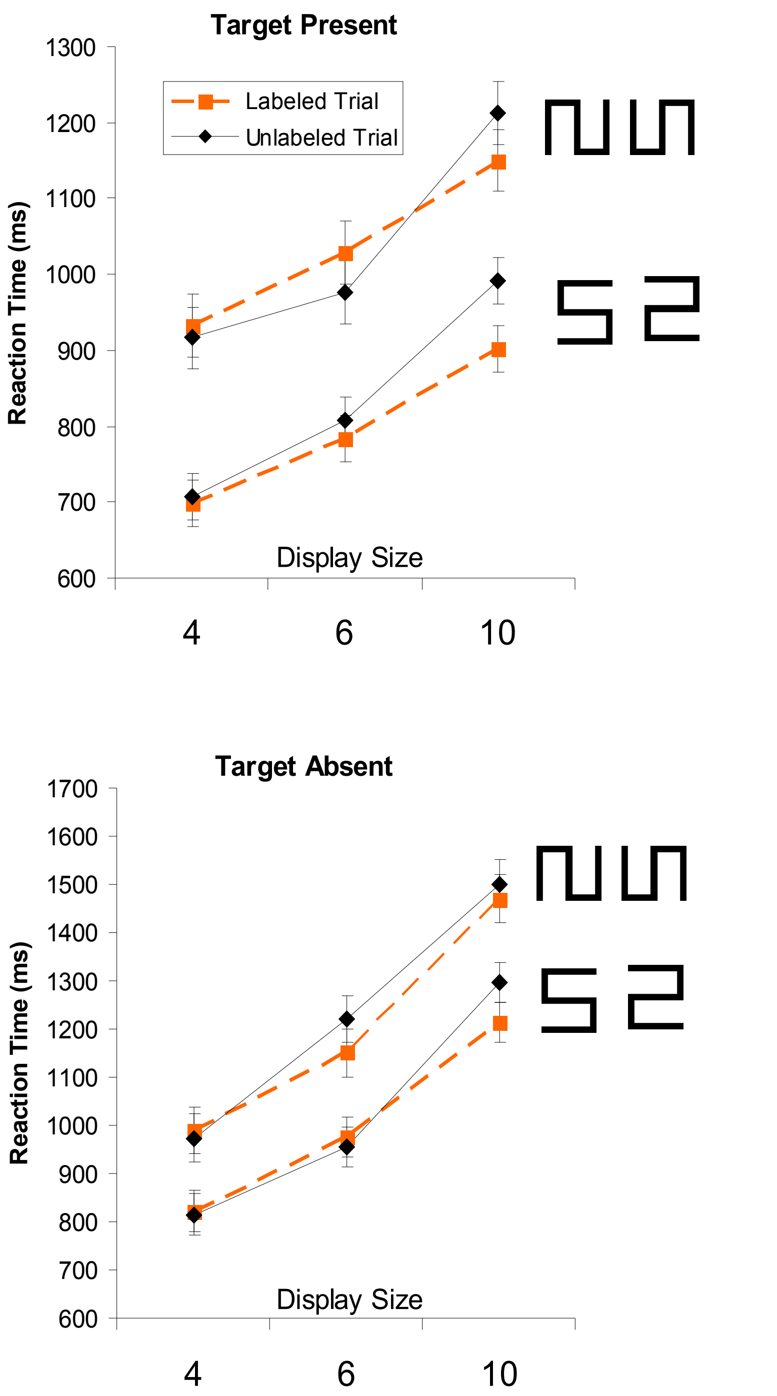

Figure 8: Experiment 1b search times for labeled vs. unlabeled upright and rotated trials. Labels name the non-targets. Labeling and orientation are intermixed (manipulated within each block). Error bars are 95% Within-Subject CIs.

There was no overall effect of labeling on RTs, $F(1, 23) = 2.02$, $p > .16$. Search slopes were reduced for labeled trials as revealed by a highly significant labeling × display-size interaction, $F(2, 23) = 5.76$, $p < .01$. There was also a significant orientation × labeling interaction, $F(1, 23) = 6.52$, $p < .025$ suggesting that the effect of labels was mediated by orientation. Analyzing the upright and rotated trials separately clarified the effect of labels. For the upright trials, hearing the distractors labeled with their category resulted in both faster overall search, $F(1, 23) = 8.1$, $p < .01$, and more efficient search (i.e., shallower slopes), $F(2, 23) = 3.27$, $p < .05$. On rotated trials, labels did not reduce overall search $F(1, 23) < 1$, but again produced shallower slopes, $F(2, 23) =$

4.07, $p < .025$ (Figure 8). It therefore appears that labels had a larger facilitating effect on upright compared to rotated trials. The target-absent trials mirrored these orientation × labeling interactions. Search was much slower, but not less efficient on rotated trials, and labeling produced more efficient search only on upright trials, $F(2, 23) = 3.63$, $p < .05$.

Table 1: Search slopes (ms/item) for Experiments 1a-b for target-present and target-absent trials with and without labels (all trials intermixed).

Experiment	Label Condition	Target Present	Target Absent
Experiment 1a	"find the…"	ms/item	ms/item
Upright Trials	Labeled	37	70
	Not Labeled	39	70
Rotated Trials	Labeled	51	95
	Not Labeled	53	83
Experiment 1b	"ignore…"		
Upright Trials	Labeled	34	65
	Not Labeled	47	81
Rotated Trials	Labeled	35	80
	Not Labeled	48	85

4.6.4 Discussion

Although participants always knew the identity of targets and non-targets, hearing the non-targets labeled facilitated perceptual processing, reducing overall RTs and resulted in shallower search slopes (i.e., more efficient search). This finding provides evidence against the possibility that labels reduce search times simply by drawing attention towards the named item. Note that such a finding is itself novel—attention has been shown to be modulated by uninformative spatial terms like "above" (Logan, 1995; Hommel et al., 2001)—but never before by uninformative object labels.

As discussed in Section 3, the Label Feedback Hypothesis predicts that labels modulate perceptual processing through top-down feedback, facilitating the processing of the display elements that possess visual features typical of the named category (i.e., elements whose features are more strongly associated with the label). On this account, hearing the label shortcuts visual processing by "pre-empting" top-down feedback, which, when compatible with the bottom-up activity from the visual display may result in shorter or fewer competition cycles needed to discriminate the target from the non-targets (Reali et al., 2006; Spivey & Dale, 2004).

Although labels decreased search slopes similarly for upright and rotated trials, the reduction of slope was accompanied by faster RTs only for the more typical / familiar upright stimuli, suggesting that the effect of labels is sensitive to the visual properties of the stimulus (this idea is explored in more depth in the subsequent experiments). This finding is compatible with earlier results showing that labeling affected recognition memory for typical, but not atypical items. According to the LFH, this interaction arises because the atypical exemplars have weaker associations to the label (for instance, it takes longer to classify them—see Experiment 3b). The top-down feedback induced by hearing "two" is most effective when items in the display match the features associated with the cue (e.g., Wolfe et al., 2004).

4.7 *Experiment 2: "Rotated two" vs. ⊓⊔ : The role of meaningfulness in visual processing*

One curious finding from Experiments 1a and 1b was the lack of a reliable difference between search slopes between the familiar (upright) and unfamiliar (rotated) stimuli. This is puzzling both because search through perceptually unfamiliar non-targets is known to be much more inefficient than search through perceptually familiar non-targets (Frith, 1974; Richards & Reicher, 1978; Shen & Reingold, 2007) and because a study using the exact stimuli as used here, revealed a more than twofold increase in search slopes for unfamiliar (80 ms/item) compared to familiar (30 ms/item) 2's and 5's (Wang, Cavanagh, & Green, 1994).

What explains the discrepancy? In studies investigating effects of familiarity, perceptual familiarity is confounded with meaningfulness—stimuli that are perceptually novel are also

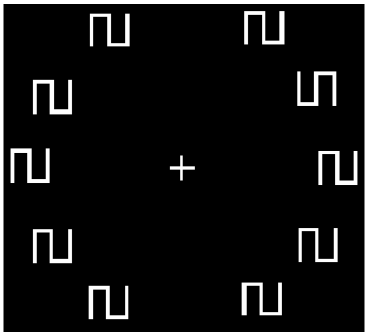

Figure 9: A sample trial from Experiment 2.

meaningless, and stimuli that are familiar have meaning (an N is known to be a letter, to have a certain sound, etc.). Thus, poor performance on unfamiliar stimuli may be due to a failure to represent them as meaningful category members rather than simply a difference in experience[14]. If so, ascribing meaning to otherwise unfamiliar stimuli should facilitate perceptual processing. This was the goal of the present experiment.

4.7.1 Participants

Sixty-one subjects, 18-22 years old volunteered for the experiment in exchange for course credit or $7. None of them had previously participated in any visual search experiments with similar stimuli.

4.7.2 Stimuli and Procedure

The stimuli were the symbols ⊔⊓ and ⊓⊔. Each trial proceeded exactly as in Experiments 1a-1b except there were no auditory labels and only the two stimuli above were used. Participants were randomly assigned to one of two groups. Participants in the *number* group were told to think of the targets/distractors as rotated 2s and 5s. This instruction was omitted for participants in the *symbol* group. In one part, participants were instructed to find a ⊓⊔ among ⊔⊓. During the other part, the target and distractor identity was reversed (with the order counterbalanced between participants). At the start of each part the target was shown on the screen accompanied

[14] The most well known familiarity effect in visual search is the finding that finding a И among N's is easy, while finding a N among И's is very difficult (e.g., N) (Frith, 1974; Wolfe, 2001). This asymmetry disappears in bilingual participants familiar with the Romance N and the Cyrillic И (Malinowski & Hubner, 2001), but notice that for the bilinguals it is not only that both N and И are perceptually familiar, they are also meaningful, representing two different letters. Thus meaningfulness and familiarity are again confounded.

with a reminder that it should be viewed as a rotated number (*number* condition). For the *symbol* group, the target was shown by itself.

Each part consisted of 10 blocks of 6 trials (target-present vs. target-absent × 3 display sizes—4, 6, or 10 × two replications). Trial order was random with the target present on exactly half the trials. Participants gave 2-alternative target *present* / *absent* responses using a gamepad controller. Participants were instructed to respond as quickly as possible without compromising accuracy. If the accuracy dipped below 92% for 24 trials, participants saw a display asking them to try to be more accurate. Response mapping (left hand *present* vs. right-hand *present*) was counterbalanced between participants. Each part began with 12 practice trials. The inter-trial interval was 750 ms. Feedback in the form of a buzzing sound was provided for incorrect responses. A sample search trial is shown in Figure 9.

After completing the experiment, participants were given a written questionnaire that asked whether they thought of the symbols ⊓⊔ / ⊔⊓ as any kind of number or letter, and if so which one(s). Participants were also queried about their consistency of label use. A question asked, "did you use this label at the beginning / middle / end of the experiment (circle all that apply)." Consistency was coded based on whether participants claimed to use labels throughout the task (*consistent group)*, or only for part of the task (*inconsistent group)*.

The questionnaire was necessary because participants in the *symbol* condition may have considered the stimuli as meaningful on their own without external experimenter-provided instructions. Conversely, participants in the *number* condition may have failed to conceive of the stimuli as rotated numbers despite the instruction to do so.

4.7.3 Results

For the *symbol* group, the responses fell into three categories. First, participants who consistently self-labeled the stimuli, either as rotated 2s/5s, or thought of them as other (often creative) symbols/symbol combinations (N=14). For instance, several participants thought of ⊓⊔ / ⊔⊓ as NU / UN, respectively. Second, participants who labeled the stimuli inconsistently (i.e., only part of the time) (N=11). Third, participants who did not report labeling the stimuli (N=16). Participants in the *number* condition fell into two categories: those who reported consistently thinking of the stimuli as rotated 2s and 5s, as instructed (N=15), and those who although instructed to do so, did not report labeling the stimuli (N=5).

The mean proportion of misses was 8% and did not differ among conditions, $F(1,60)<1$. The false alarm rate, however, was greater in the *symbol* condition (M = .05) than the *number* condition (M = .01), $F(1,60) = 4.74$, $p < .05$. Reaction time (RT) analyses that follow include only correct responses. Responses with RTs shorter than 150 ms. and RTs greater than 3 standard deviations of participants' means were excluded from the analyses. Analyses were conducted using ANOVAs with display size as a within-subjects factor, and instruction-condition as a between-subject factor. Group RTs are shown in Figure 10 for target-present trials (top) and target-absent trials (bottom). Results for target-present and target-absent trials closely paralleled each other, so the subsequent analyses focus on target-present trials.

Initial analysis with instruction-condition as the between-subject factor revealed significantly slower responses for the *symbol* group, $F(1,60) = 7.31$, $p < .01$. The condition × display-size interaction was not significant, $F(2, 60) < 1$. Correlating performance with questionnaire responses made it apparent that participants' judgments of whether they thought of the characters in terms of any familiar symbols predicted performance. Using the 5 groups derived from the questionnaire responses as the between-subjects variable revealed a significant main effect $F(4, 57) = 6.47$, $p < .0005$, and a significant group × display size interaction, $F(5,57) = 2.26$, $p < .05$. Participants in the symbol condition who reported consistently thinking of the characters in terms of familiar categories had mean RTs that were statistically indistinguishable from participants who were explicitly told to think of the characters as numbers, whether or not the latter reported thinking of the characters as numbers: target-present RTs for the participants assigned to the symbol group who reported consistently labeling the stimuli: M =1114 ms, SD = 199 ms, versus for participants assigned to the *number* condition, M =1143 ms, SD=237 ms., $F(2,32) < 1$. These three groups were therefore collapsed into a single group for the subsequent analysis (see Figure 10). Participants in the *symbol* group who reported either inconsistently labeling the stimuli, or not labeling them had mean search times were not significantly different, $F(1,25) < 1$. The search slopes also did not differ, $F(2,25) = 1.54$, $p > 2$. These two groups were therefore also collapsed into a single group. A mixed ANOVA with display-size as a within-subjects factor and the two collapsed groups as a between-subjects factor revealed a highly significant difference in mean search times, $F(1,60) = 26.36$, $p < .0005$, and a significant group × display-size interaction—47 ms/item versus 66ms/item for target-present trials), $F(2, 60) = 3.77$, $p < .05$, and 110 ms/item versus 153 ms/item for target-absent trials, $F(1, 60) = 8.81$, $p < .0005$.

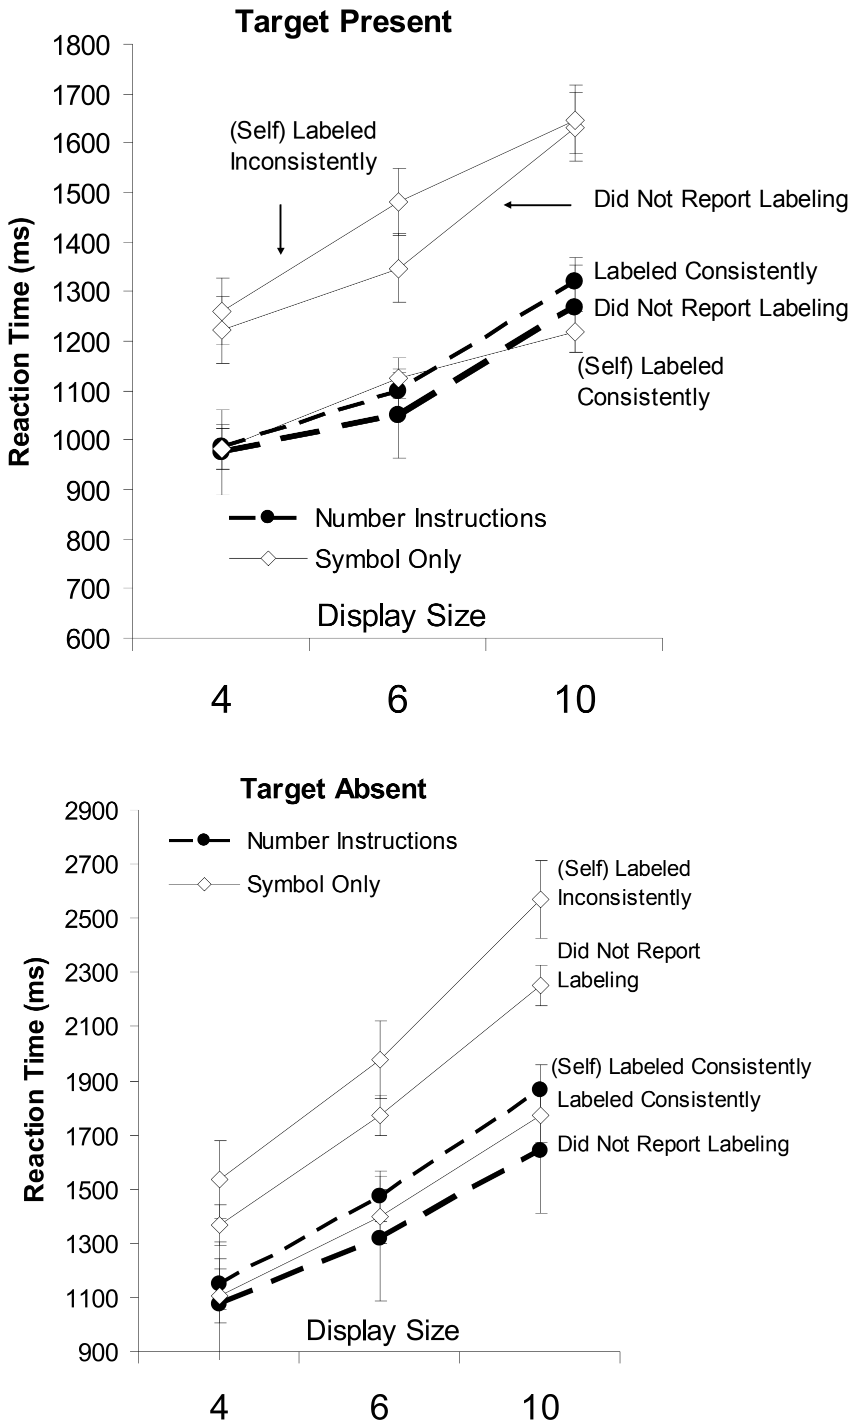

Figure 10: Search performance in Experiment 2 as a function of instruction and reported use of labels.

Figure 11 shows performance over the course of the experiment, split by the original two randomly assigned conditions (*number* vs. *symbol*) (top), split into groups derived from questionnaire response (middle), and re-collapsed into two new groups—collapsing participants who were instructed think of stimuli as meaningful or who self-labeled stimuli consistently vs. those assigned to the *symbol* condition (bottom). Several patterns are apparent: First, the *number* group starts showing an RT advantage early on in the experiment. The block × condition interaction was not significant when tested with the original *number* and *symbol* groups, F=1.04, but becomes significant when self-labeling performance is taken into account (middle), $F(19, 884) = 1.48$, $p = .006$. The overall pattern is seen more easily in the bottom figure. Individuals assigned to the *number* condition that did not consistently self-label the stimuli show no improvement over the course of the experiment, while participants who were told to think of the items as rotated numbers, or who consistently ascribed meaning to the perceptually novel items, showed improvement as indicated by a significant block × condition interaction, $F(19, 947) = 1.61$, $p < .05$[15].

4.7.4 Discussion

[15] The small number of observations in each cell resulted in a number of statistical outliers (Std. residuals > 2.95) which were iteratively removed from the analyses.

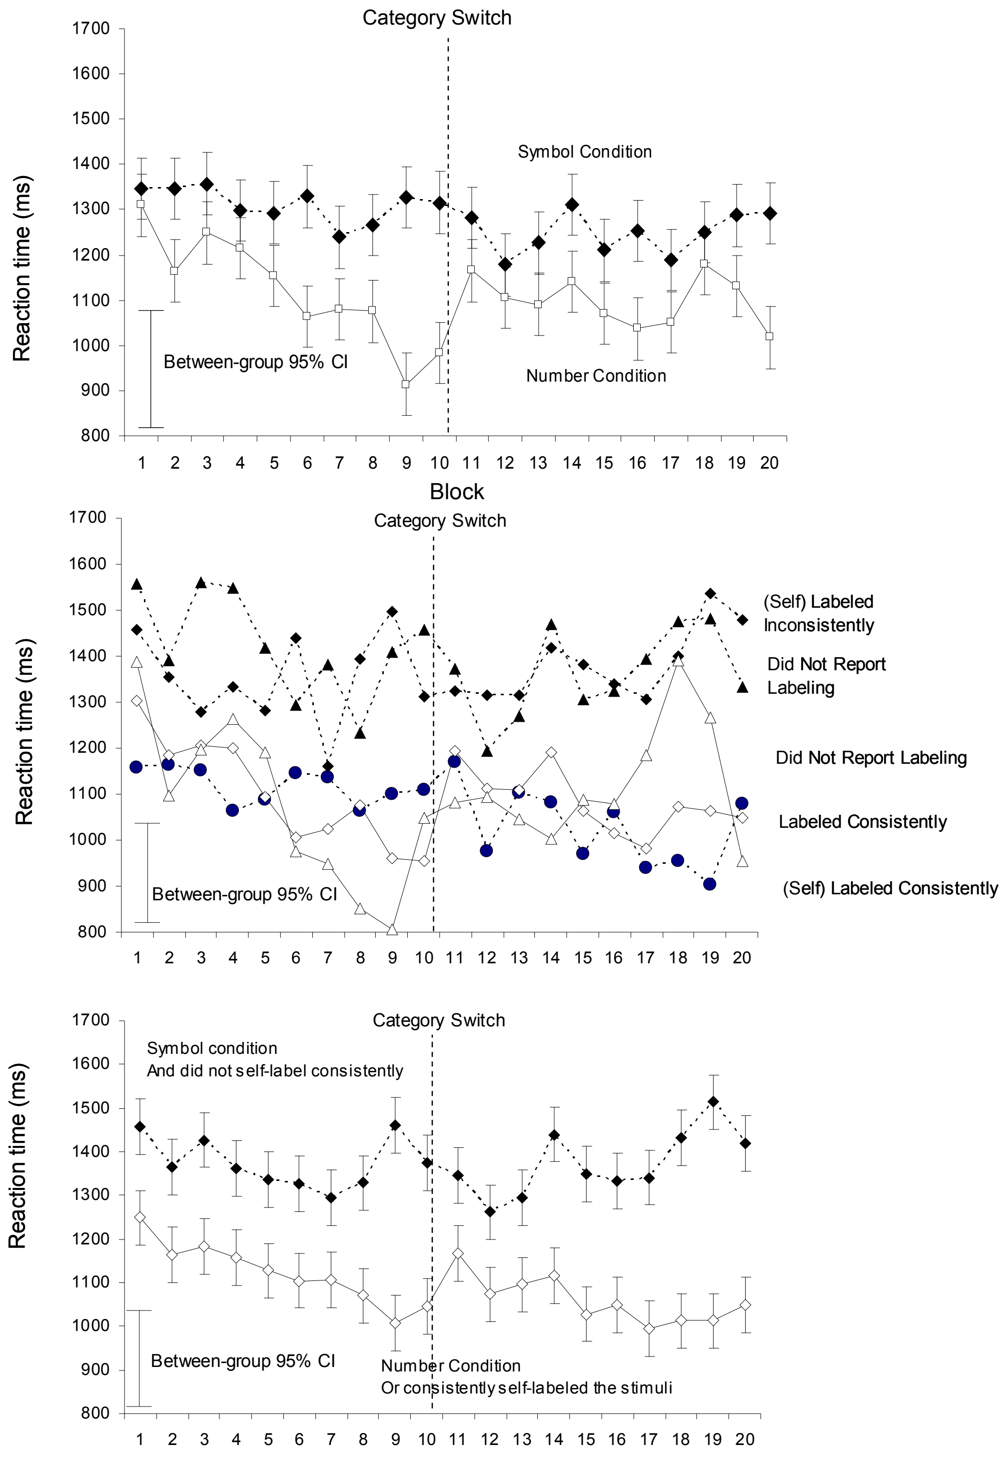

Figure 11: Search performance in Experiment 2 as a function of block. The top panel shows the original randomly-assigned groups. The middle panel shows performance separated by group membership and category responses (filled = Symbol Condition, unfilled = Number Condition). The bottom panel shows performance collapsed by instructed or consistently self-produced labeling (see text).

Instructing participants to consider novel stimuli as instances of a familiar category significantly improved mean search times and search efficiency. The benefit of representing perceptually novel items as members of meaningful categories was also observed in individuals who reported *consistently* self-labeling the stimuli without being told to do so. Participants who treated the stimuli as meaningless continue to show inferior performance even after hundreds of trials, showing no appreciable practice effects. In contrast participants who were either assigned to the *number* condition or consistently self-labeled the stimuli, not only show superior performance from early in the experiment, but show further improvement during the task.

These results suggest that theories of visual processing that do not take meaningfulness into account (Duncan & Humphreys, 1989; Pylyshyn, 1999; Theeuwes, 1993; Treisman & Gelade, 1980) may need to be reassessed. The argument that perceptual processing in visual search depends on more than visual similarity has been recently made by Rauschenberger and Yantis (2006) who argued that perceptual encoding depends on stimulus *redundancy*. For instance, not all combinations of a circle and a line are equally redundant. Combinations that create the letter Q are processed more efficiently. A Q, being a member of an implicit set of size 1, is highly redundant compared to the harder-to-process  —a member of a less redundant set 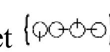 (Garner & Clement, 1963). Insofar as the redundancy framework is a useful one, it is clear that redundancy cannot be reduced to visual features, but must take meaningfulness into account. Experiment 2 shows that controlling for all perceptual variables, meaningfulness in its own right affects perceptual processing.

The interplay between bottom-up and top-down processes is further highlighted by Experiments 1a-1b. Even though upright numbers were both meaningful and familiar, on-line presentation of labels further facilitated performance on the search task, leading to shorter RTs (Experiments 1a-1b) and shallower slopes (Experiment 1a). Labels were less effective in modulating on-line performance for the rotated trials. Together, the experiments show that while search for rotated items is only slightly affected on-line by labeling, there is a robust facilitation when participants ascribe meaning to the stimuli.

4.8 Experiment 3a: Effects of Visual Quality

One way in which the labels might facilitate perceptual processing is by "cleaning up" the perceptual representations of category-typical stimuli. This would increase the homogeneity of all the labeled items perhaps facilitating search through increased perceptual grouping. If this is the case, then the facilitating effect of labels should increase under conditions of lower visual quality, to the extent that decreasing visual quality increases the potential for the cleaning up of visual representations). Alternatively, a decrease in visual quality may disrupt the link between the stimulus and its category. Thus, stimuli of lower visual quality may be less affected by category knowledge, and hence, less associated with the category names, rendering null any facilitative effect of the labels. Experiments 3a-3b attempt to discriminate between these two alternatives.

4.8.1 Participants

Twenty-one Carnegie Mellon University undergraduates volunteered for the experiment in exchange for course credit. Three participants were eliminated for having accuracy below 80%.

4.8.2 Stimuli and Procedure

The procedure was identical to Experiment 1b (section 4.6) with the following exceptions: Rather than varying display size, the present experiment manipulated visual quality. Elements in the search displays were degraded by various amounts through occlusion of randomly-placed

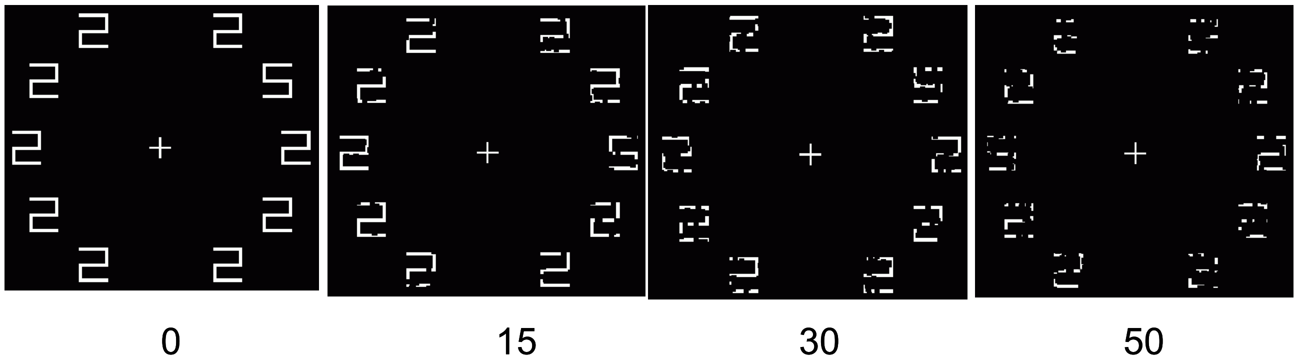

Figure 12: Examples of the four visual quality conditions used in Experiment 3a.

small squares. Examples of stimuli subjected to occlusion levels of 0, 15, 30, and 50 can be seen in Figure 12. These occlusion values refer to the number of squares occluding each character.

Each participant completed 2 counterbalanced blocks: searching for an upright/rotated "digital" 2 among upright/rotated 5's and searching for an upright/rotated 5 among upright/rotated 2's. Each block began with 12 practice trials, followed by 320 search trials, all containing 10 display elements. The verbal cues were "ignore two" or "ignore five" for the *label* condition and "ignore [noise]" for the *no-label* condition.

4.8.3 Results

Incorrect responses and those less than 150 ms or greater than 6000 ms were excluded from the analyses. Responses were entered into an ANOVA with subjects as random factors and the remaining factors as fixed: labeling (*label* versus *no-label*), orientation (upright versus rotated), and occlusion amount (0, 15, 30, 50). Analyses will concentrate on target-present trials, except where noted. Search performance is shown in Figure 13. Visual quality had a profound impact on performance with greater RTs for more visually degraded stimuli, $F(3, 20) = 45.54$, $p < .0005$. Errors also increased with more degraded stimuli, from 7% at no occlusion to 13% at the highest level of occlusion, $F(3, 20) = 34.38$, $p < .0005$. No other effects of errors were found, so the remaining analyses will focus on RTs.

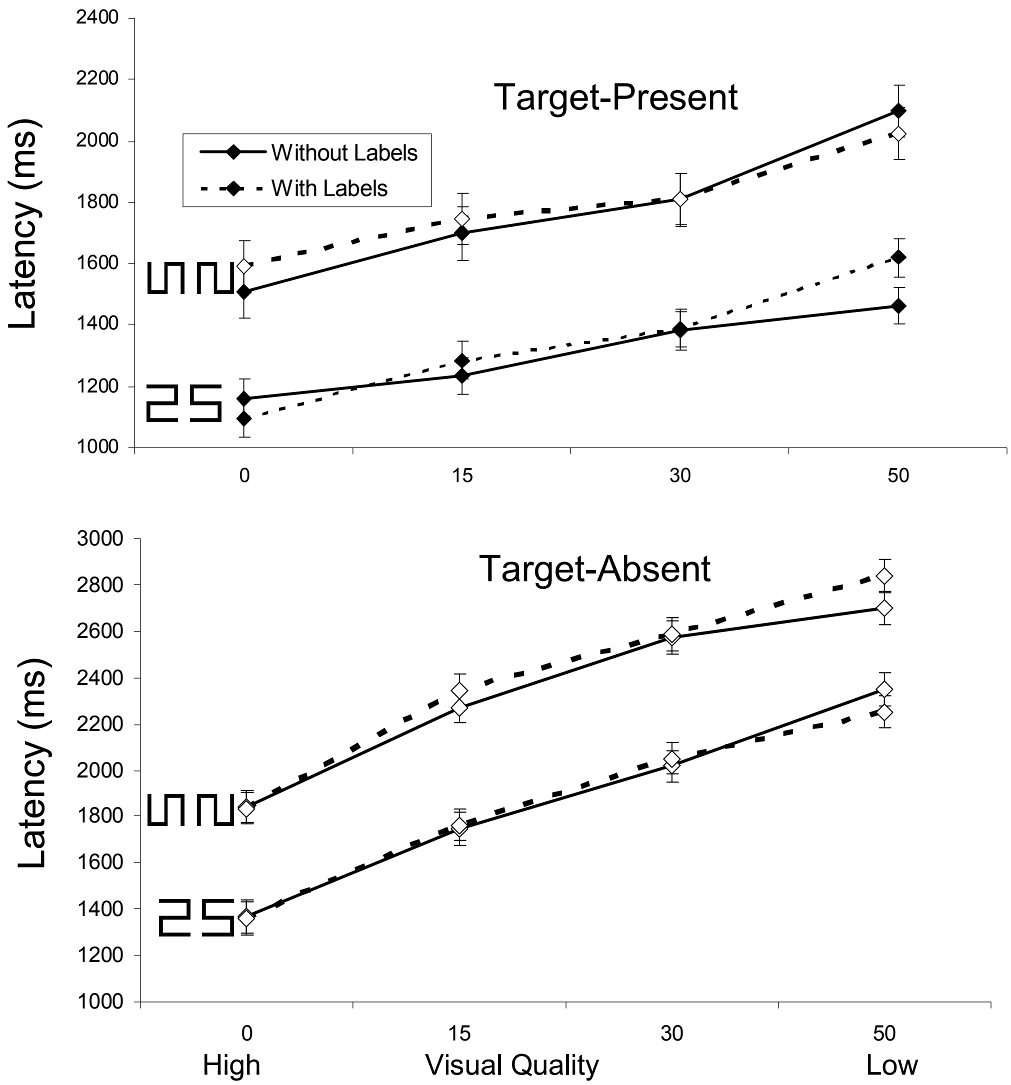

Figure 13: Search performance in Experiment 3a, as a function of orientation, visual quality, and labeling condition. Error bars are within-subject 95% CIs.

As before, search through upright numbers was faster than search through rotated numbers, $F(1, 20) = 56.48$, $p < .0005$. Search through upright numbers was also more accurate, $F(1, 20) = 32.83$, $p < .0005$. There was no interaction between visual degradation and orientation for RTs, $F(3, 60) = 1.44$, $p = .24$, or accuracy, $F(3, 60) = 1.21$, $p = .32$. Though there was no overall interaction between visual quality and orientation, a curious effect emerged when N-T category was included in the analysis. Recall that all participants searched for a 2 among 5's in one block

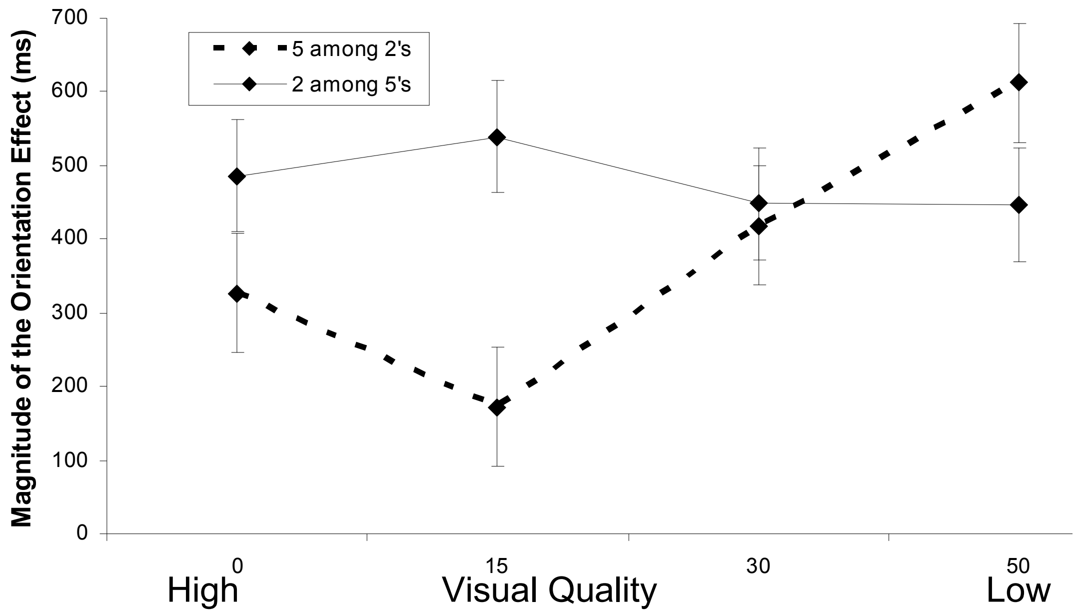

Figure 14: The category asymmetry effect from Experiment 3a. The difference in searching among upright and rotated stimuli (orientation effect) as a function of visual quality. Searching among 2s is easier than searching among 5s for high-quality stimuli, but is more difficult for low-quality stimuli. Error bars are within-subject 95% CIs.

and for a 5 among 2’s in another block. Including N-T category as a within-subject factor revealed a three-way interaction between orientation, visual quality and N-T category, $F(3, 60) = 3.88$, $p = .01$. Examining separately the orientation × visual quality interactions for the two N-T categories revealed no interaction when 5s were the non-targets, $F<1$, but a significant interaction when 2s were the non-targets, $F(3, 60) = 4.39$, $p = .007$. When searching for a 5 among 2s, rotated stimuli were more affected by decreasing visual quality than upright stimuli. In comparison, visual processing of rotated 5s survived a greater levels of visual degradation. Thus, the “orientation effect,” ($RTs_{rotated}$-$RTs_{upright}$) was not affected by visual quality when searching among 5s, but increased when searching among 5s (Figure 14). This curious finding is discussed in section 4.8.4.1 below.

There was no main effect of labeling condition, $F(1, 20) = 1.29$, $p = .27$. Importantly, there was a significant three-way interaction between labels, orientation, and visual degradation, *F*(3,

60) =5.26, p =.003. Labels had no impact on rotated trials, Fs<1, though there was a trend for labels to slow down search for the non-degraded rotated trials, $F(1, 20) = 2.85$, $p = .11$.

There was a significant interaction between label condition and visual quality for the upright trials, $F(3, 60) =4.61$, p =.006. In particular, labels significantly slowed down performance on the most degraded upright trials, $F(1, 20) =10.09$, p =.005, and marginally improved performance on the non-degraded upright trials, $F(1,20) =3.44$, p =.079. A further analysis of the latter condition revealed 2 extreme data points (one 2.3 SDs above the mean, another 2.3 SDs below the mean) which inflated the variability. A nonparametric re-analysis of the data (Wilcoxon sign-rank test) found a significant difference between the conditions (*Mdn* Difference = 62 ms), T(16) = 174.0, p = .04. A parametric re-analysis of the data removing the two extreme points likewise found significantly shorter RTs in the *label* condition, $F(1,18) = 6.00$, p = .025.

4.8.4 Discussion

Labels only facilitated search for the intact stimuli, effectively replicating the finding of Experiments 1a-1b. Even slight visual degradation of the stimuli removed the label-facilitation effect. This finding is consistent with the idea that however labels modulate visual representations, this modulation is critically dependent on an association between the visual stimuli and labels and any disruption of this association eliminates the on-line component of the label-effect. Interestingly, for the most degraded stimuli labels significantly slowed down responses. This would be expected if labels deploy a kind of category template. According to this reasoning, hearing “ignore 5” biases the visual system to process category-typical 5’s. Encountering a highly degraded 5, which is arguably difficult to categorize, produces a mismatch between the top-down activity priming induced by the label and what is actually seen (bottom-up activity), leading to increased RTs on the *label* trials. Though it is still possible that the mechanism of the label-facilitation effect might involve some type of perceptual clean-up, it seems clear that this modulation of visual representations by labels is disrupted by lowering the visual quality of the stimuli. The present experiment thus shows that labels do not facilitate performance on-line for visually degraded stimuli. It should be noted that, as in Experiments 1a-1b, participants were explicitly told that all the stimuli were instances of the numbers 2 and 5. The lack of a facilitating effect of labels for the degraded stimuli is thus in contrast to trials in which participants are arguably attempting to classify the stimulus as a 2 or 5, but do not hear it labeled. If one were to contrast the labeling condition with a condition in which participants are not told that the degraded stimuli are 2s and 5s, it is entirely plausible that performance on the degraded stimuli would be superior in the label condition. However, such an experiment would confound the on-line effect of labels with the role of meaningfulness. The purpose of the present experiment was to explore the on-line role of labels while controlling for such factors as meaningfulness.

An explanation of the present results in terms of the LFH is that search RTs are closely tied to how easily classified a stimulus is into a meaningful category (here, 2's and 5's). Recall that in Experiment 2, search through identical stimuli was strongly determined by whether participants were told to think of the stimuli as rotated numbers. Insofar as visual degradation may impair classification, it may impair search. Experiment 3b tests this hypothesis.

4.8.4.1 A Category Asymmetry Effect: A relationship between frequency, category tuning and label-facilitation

The present experiment found that effects of visual quality were mediated by stimulus category: visual degradation slowed search more when searching for a rotated 5 among rotated

2s compared to a rotated 5 among rotated 2's. That is, with increased perceptual processing demands, the orientation effect ($RTs_{rotated}$-$RTs_{upright}$) grew larger for 2's but not 5's (Figure 14).

How can this be explained? Because the 2 and 5 stimuli are just mirror images of each other, perceptual confounds can be safely ignored. It is also clear that 2s and 5s are highly overlearned categories. But in fact, there is a large frequency difference in the rate of encountering both the words "two" and "five" and their corresponding numerals. According to the British National Corpus, the frequency of the word "five" is 40,738 / 100 million. The word "two" is almost four times more frequent at 156,111 / 100 million. In written texts, the numeral "2" occurred 34,394 times while the numeral "5" occurred 17,557 times, the great majority of them, presumably as normally-oriented typical category exemplars. Greater exemplar frequencies may establish a stronger central prototype, making the 2 category in effect narrower or sharper than the 5 category. This hypothesis leads to two predictions: First, a logical consequence of this difference in the strength of the category attractors may be faster processing for the typical stimuli, but at the cost of slower processing for the more atypical/degraded stimuli. This is exactly what is seen in Figure 14. Second, category strength may be closely linked to effects of labels on perceptual processing—a stronger central prototype should elicit a greater modulation of visual representations by category labels. One test of this prediction is to use the magnitude of the orientation effect as a proxy for category tuning with a smaller orientation effect arising from a less selective, i.e., more broadly tuned category. We can then compute a correlation (Pearson product-moment coefficient) between the magnitude of the orientation effect for individual participants and the observed facilitation due to labels (for upright, non-degraded stimuli) for these same participants. This correlation was positive and significant, $r(19) = .438$, $p < .05$, indicating that an increased orientation effect (more narrow category tuning) predicts a greater label advantage.

4.9 Experiment 3b: The Classification-Discrimination Link

Does degrading the stimuli really make them more difficult to classify? After all, stimuli quite significantly degraded (e.g., level 30 in Figure 12) seem to be easily identifiable as 2s and 5s. To find out, a new group of participants classified the individual stimuli used in Experiment 3a. Then, to determine the degree to which classification RTs predicted search RTs, the classification RTs were correlated with the search RTs from Experiment 3a. A high correlation between classification RTs and visual search RTs provides further evidence for the role of categories in visual processing.

4.9.1 Participants

Ten Carnegie Mellon University undergraduates volunteered for the experiment in exchange for course credit.

4.9.2 Stimuli and Procedure

Participants completed 320 classification trials in which they were asked to classify upright and rotated 2s and 5s at various levels of visual quality, by responding with one key for “2” and another for “5.”. To ensure high accuracy and stable performance, the first 64 trials included accuracy feedback—a bleep sound following correct responses and a buzzing sound following incorrect responses. These feedback trials are not included in the subsequent analyses (though their inclusion does not qualitatively change any of the reported analyses). To prevent anticipatory responses, the ISI varied randomly between 500 and 900 ms.

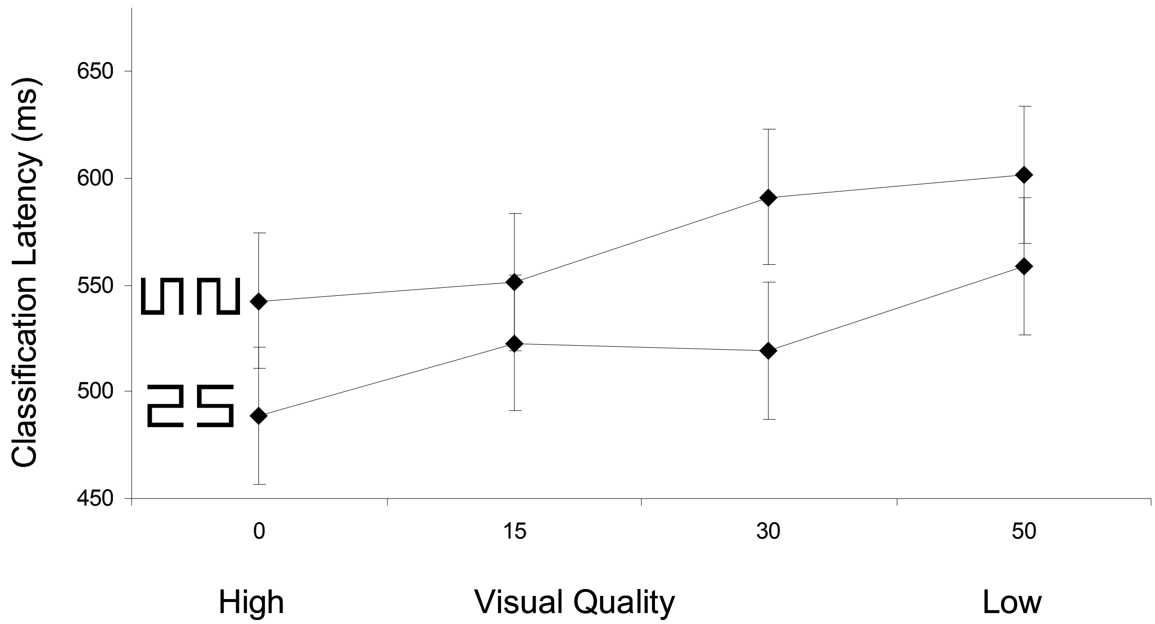

Figure 15: Classification performance in Experiment 3b as a function of orientation and visual quality. Errors bars are within-subject 95% CIs.

4.9.3 Results

Incorrect responses and those less than 150 ms or greater than 1500 ms were excluded from the analyses. Upright stimuli were classified more quickly than rotated stimuli, $F(1, 9) = 8.09$, $p = .02$ Classification RTs correlated with visual quality, $F(3, 27) = 4.77$, $p < .01$. There was no interaction between orientation and visual quality condition, F<1 (Figure 15). Accuracy was high (98% for upright; 97% for rotated) and did not correlate with orientation or visual quality.

To determine whether classification RTs predicted search performance, the mean RTs at each level of visual quality were entered into a regression analysis as a predictor for search RTs at the identical visual qualities. The results are shown in Figure 16, with separate regressions for upright and rotated stimuli. It is immediately apparent that visual quality and classification RTs are very good predictors of search RTs. In fact, the Pearson correlation between visual quality and search RTs for upright trials is 1.0. It is also apparent that even collapsing across the orientation of the stimulus, classification RTs faithfully predict search RTs for all visual

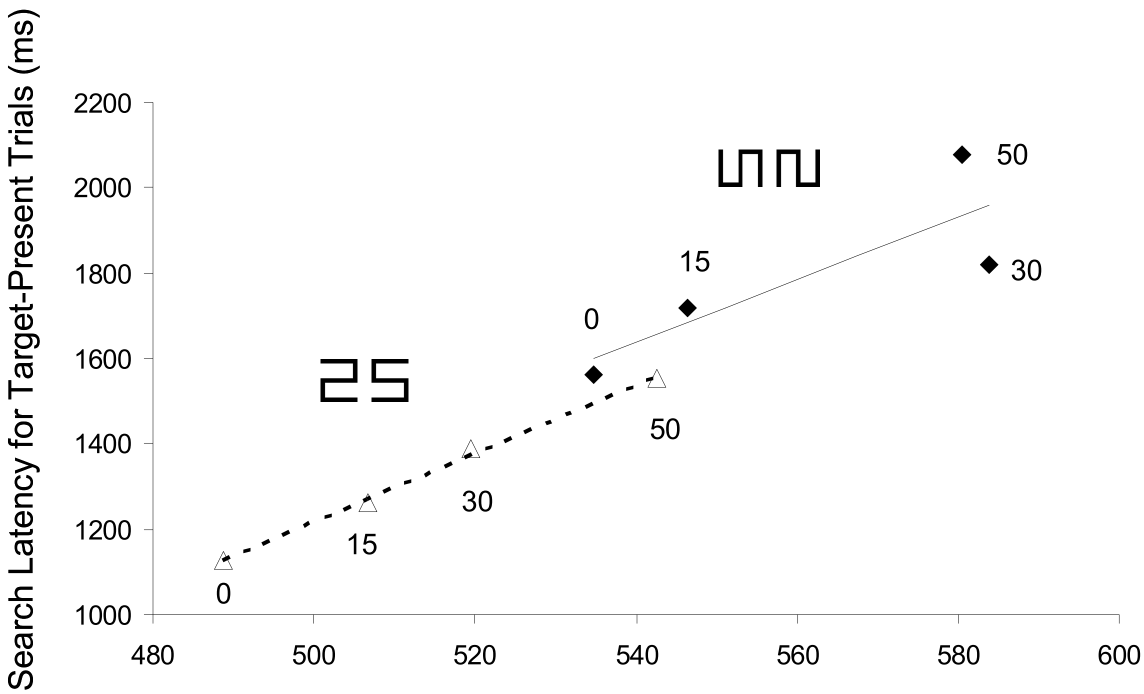

Figure 16: The relationship between search RTs (Experiment 3a) and classification RTs (Experiment 3b) for target-present trials for stimuli of different visual quality and orientation. The overall correlation is .961.

qualities, $r(6) = .961$, $p < .0005$. The slope is 8.89, i.e., each extra ms spent classifying a single character predicted an increase of approximately 9 ms in searching for a target among 9 non-targets.[16]

The finding of a category asymmetry in Experiment 3a combined with the present finding of an extremely tight relationship between classification and search RTs leads to a prediction that the category asymmetry effects observed in search (Experiment 3a) should be mirrored in classification RTs in the present experiment. Specifically, participants may be faster in classifying typical 2s than typical 5s, but faster in classifying atypical 5's than atypical 2's. This prediction was tested by adding category (2 or 5) as a within-subject factor into the ANOVA analysis performed above. Category did not interact with visual quality, F<1, but interacted marginally with orientation, $F(1, 9) = 4.59$, $p = .06$, shown in Figure 17. The interaction was in the predicted direction. Classification RTs for atypical (rotated) 5s were indeed shorter than for

[16] The full regression equation is search RT = -3218 + 8.89 * classification RT

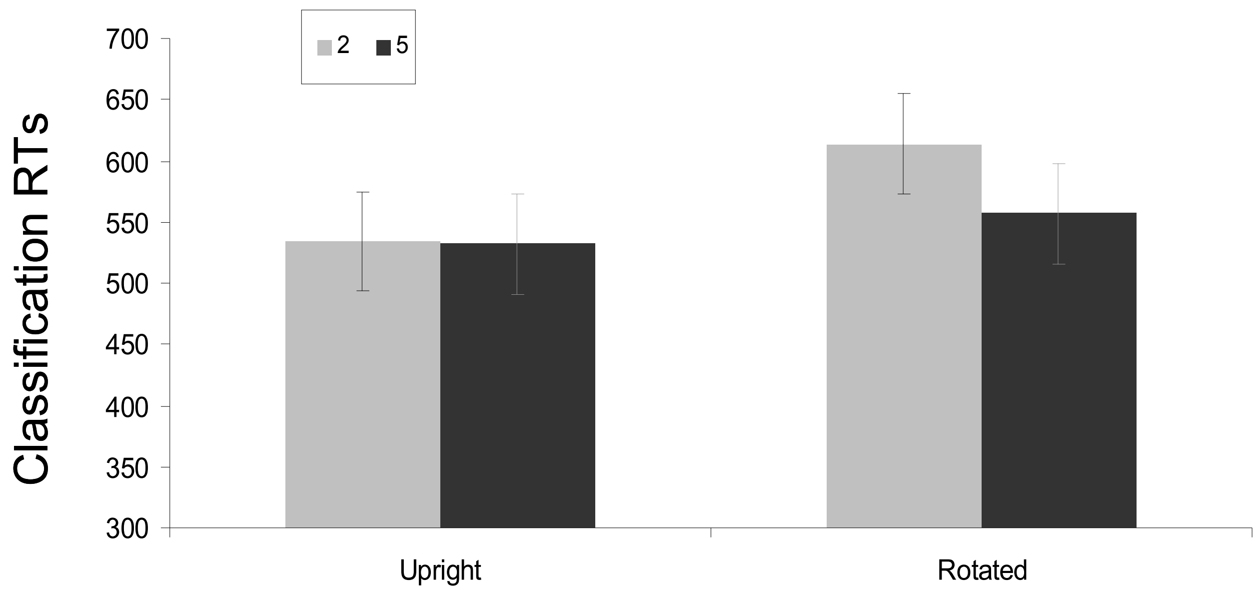

Figure 17: Classification Performance (Experiment 3b). Rotated 2s are harder to classify than rotated 5s. A likely cause of this difference is the greater ecological frequency of 2s compared to 5s (see text).

rotated 2's, $F(1, 9) = 3.67$, $p < .05$, supporting the hypothesis that the "2" category is more narrowly tuned than the "5" category, possibly owing to a difference in frequency of encounter[17].

4.9.4 Discussion of Experiments 3a-3b

A facilitation due to labels was observed only for the non-degraded stimuli—those easiest to find. The interpretation of this result that is consistent with the LFH is that labels facilitate processing most for stimuli that are easiest to classify[18]. Experiment 3b tested the link between classification, search, and labels more directly by collecting classification RTs for the stimuli used in Experiment 3a and measuring their correlations with search RTs. The extremely tight correlation (Figure 16) highlights the intimate connection between stimulus classification and perceptual processing. Lowering visual quality made both search and classification more

[17] It is possible that the 'digital' 5 was a better example of a the more typically rendered 5 than then digital 2 was of a typically rendered 2. One would therefore predict that classification for upright 5s would be faster than for the upright 2s, which was not observed (Figure 17).
[18] This explanation does not rule a mechanism by benefits of hearing a label results from the grouping together of the labeled items. The experiments presented in the Appendix provide demonstrations of "conceptual grouping" at work.

difficult, and as classification became more difficult, labels started to impair search. As shown in Figure 16, overall search performance is best predicted by the time it takes to classify individual stimuli. Conditions in which the target and non-target are most easily classified, lead to fastest search times. Thus, as predicted by the LFH, the effect of labels on visual processing seems to depend on the existence of a strong association between the visual stimulus and the label.

A further finding was that the difference in search times between upright and rotated stimuli (the orientation-effect) correlated with the degree to which labels decreased RTs for finding the high-quality upright stimuli. One explanation is that the orientation effect serves as a proxy measure for broadness of category tuning. Individuals with more narrowly tuned category (as measured by a large orientation effect), benefit more from hearing labels due to the more focused category representations possessed by these participants. Experiments 4a and 4b provide further tests of this explanation.

4.10 Experiment 4a: Category goodness and the label-facilitation effect

Experiment 4a had four main aims. The first was to generalize the effect of labels on visual processing to new visual and auditory stimuli. In Experiments 1a, 1b, and 3a, the audio cue in the *no-label* condition used white noise in place of the target or non-target label. Although the specificity of the observed effects of the labels make it clear that the difference between the cue conditions is specific to the labels, it is possible that hearing “find [noise]” or “ignore the [noise]” was somehow odd and contributed to the slowing of RTs on the *no-label* trials. Experiment 4a makes use of novel stimuli and replaces the “ignore the [noise]” cue with “ignore distractors.”[19]

[19] Experiments in Appendix 1 provide an additional demonstration of the labeling effect with an additional set of stimuli and labels.

One curious finding of Experiment 3a was that even small disturbances to the visual integrity of a familiar stimulus were enough to nullify any facilitation effect of labels. The second aim of Experiment 4a was to examine this effect further by smoothly varying the category goodness of the non-targets by morphing them between two categories.

Experiment 3a also showed that the orientation effect—a proxy measure of hypothesized category-tuning—was related to the strength of the label-facilitation effect. Narrowed category tuning predicted a greater effect of the labels. The third aim of this study was to examine this further by computing two new measures of category-tuning and correlating them with the influence of labels.

The fourth aim of this experiment was to obtain classification and search RTs from the same individuals. This makes it possible to further investigate individual differences. For instance, it may turn out that individuals vary on the pre-existing strength of their prototypes for one of categories, and this predicts search times and effects of category labels. Of course, using the same individuals for the search and classification tasks is potentially problematic because performing hundreds of search trials may affect subsequent classification performance, and completing a classification task may affect subsequent search performance. One solution is to manipulate the order of the search and classification tasks, thus allowing to . In Experiment 4a a classification session follows search. In Experiment 4b, search follows classification. This manipulation also makes it possible to directly investigate the effects of search trials on classification and vice-versa.

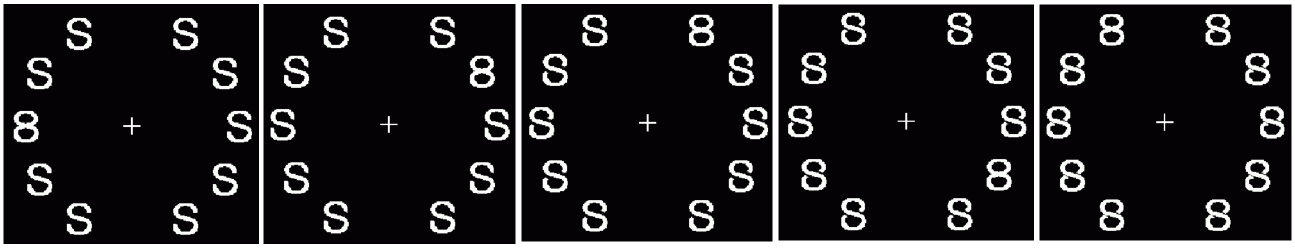

Figure 18: Five trial types used in Experiments 4a and 4b. The target is always the 8. The display viewed by participants was larger than depicted here, thus the details of the stimuli were easier to make out.

4.10.1 Participants

Twelve Carnegie Mellon University undergraduates volunteered for the experiment in exchange for course credit. Classification data from one participant was excluded due to chance-level accuracy.

4.10.2 Stimuli and Procedure

Participants searched for an 8 among S's. Rather than varying visual quality, different trials varied the perceptual distance between target and non-targets by morphing the non-target (S) towards the target (8), creating 5 levels of discriminability (Figure 18). For the highest level of discriminability (easiest trials), both the target and the non-target are readily recognized as typical category exemplars. As target/non-target similarity increases, the non-targets look less like S's and more like 8's leading to increasingly difficult search. Prior to beginning the search trials, participants were shown the target and the range of non-targets. The instructions emphasized that the target had no gaps in its outline (i.e., is a complete figure 8), while all the non-targets had gaps of varying degrees.

Each participant completed 12 practice trials, followed by 400 search trials (20 repetitions of target-present/target-absent × 5 levels of discriminability × 2 labeling conditions). All trials contained 10 display elements. On a random half of the trials, the non-targets were labeled:

“ignore S” (*label* condition) and on the remaining trials, participants heard “ignore distractors” (*no-label* condition). The cues had respective durations of 1327 ms and 1500 ms. The interval between cue offset and the appearance of the search display was equated such that the display always appeared 700 ms after cue offset.

Following the search trials, participants completed a session of 180 classification trials in which they were shown stimuli previously seen in the search session, and asked to classify them as 8’s or S’s as quickly and accurately as possible. As in the search session, there was a single type of 8, and five variants of S’s. Because the aim of this task was to obtain RT data for classifying the S’s of varying similarities to the 8, the stimuli were presented with a 5-to-1 ratio of S’s to 8’s. No accuracy feedback was provided. To prevent anticipatory responses, the ISI varied randomly between 500 and 900 ms.

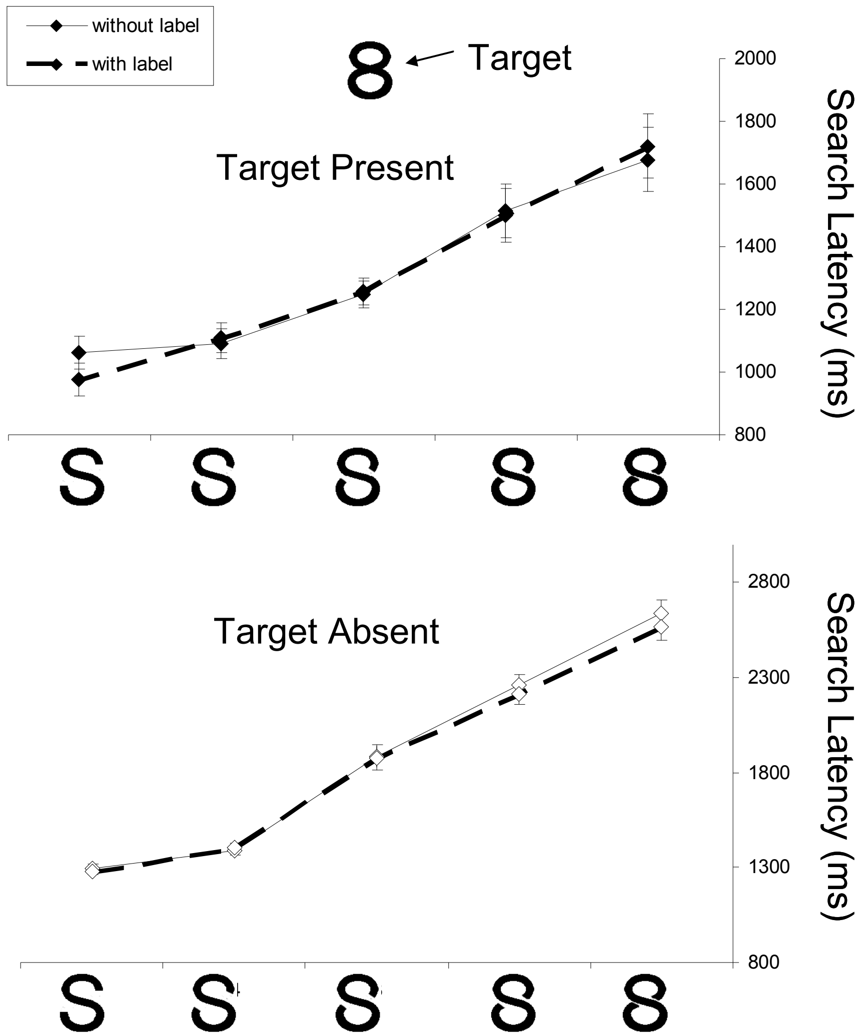

Figure 19: Search performance in Experiment 4a. The x-axis shows performance for non-targets of different degrees similarity to the target. Errors bars are within-subject 95% CI.

4.10.3 Results

4.10.3.1 Search

Incorrect responses and those less than 150 ms or greater than 4000 ms were excluded from the analyses. The pattern of search times are shown in Figure 19. Analyses will focus on target-present trials. Non-target (N-T) type and labeling condition were entered into a repeated-measures ANOVA, revealing a highly significant effect of stimulus-type, $F(4, 44) = 55.46$, $p < .0005$. The effect of labels and the label-condition × N-T interaction were not significant, F's<1.

These RT effects were mirrored by an analysis of errors. Errors increased from 9% on the easiest search trials to 25% on the hardest search trials, $F(4, 44) = 16.52$, $p < .0005$. There were no main effects or interactions with labeling condition, F's<1. An analysis of target-absent trials revealed identical effects of trial difficulty (Figure 19 bottom), and no effects of labels, F's<1.

Planned comparisons of the effects of labels on RTs for each stimulus type revealed that labels significantly affected search only for the easiest search trials, $F(1,11) = 5.90$, $p = .033$, resulting in a mean facilitation due to labels of 90 ms. There was no accuracy difference between *label* and *no-label* search trials $F(1,11) = 1.22$, $p = .294$. In fact, correlating the label RT advantage for each participants with an accuracy difference between the *label* and *no-label* trials revealed a marginal positive correlation, $r(9) = .491$, $p = .105$. Individuals who showed a smaller or nonexistent label-advantage, tended to perform more accurately on the *no-label* trials—evidence that the overall facilitation effect of labels was not due to a speed-accuracy tradeoff.

The orientation effect in Experiment 3a was computed by subtracting the RTs for the easier (upright) search trials from the more difficult (rotated) search trials. Here, an analogous measure was computed by subtracting the trials with the most typical S non-targets from the least typical S's (i.e., the trials in which the T and N-T's look most alike). A greater between the easiest and hardest search trials may indicate a more narrowly tuned category representation. To determine whether degree of category tuning was related to the effects of labels, a *label-advantage* score was computed for each participant by subtracting their RT_{label} from $RT_{no\text{-}label}$: the more positive the score, the greater the facilitation effect of the labels. Correlating these two measures revealed a highly negative correlation, $r(10) = -.719$, $p = .009$. Restricting the analysis to the easiest trials only (because it is on these trials that the effect of labels was observed), increased the correlation to -.832, $p = .001$. Thus, individuals who performed best on the easiest trials compared to the most difficult trials were the ones who showed the greatest benefit of labels. Insofar as the RT difference between easiest and hardest trials acts as a proxy for category tuning, the individuals with the most tightly-tuned category representations were the ones who most benefited by hearing the labels.

4.10.3.2 Classification

Stimulus type was a significant predictor of RTs and errors when entered into a within-subjects ANOVA, $F(5, 50) = 20.61$, $p < .0005$ for RTs, $F(5, 50) = 6.02$, $p < .0005$ for errors which increased from approximately 2% to 9%. These relatively low error rates indicate that participants understood the task (correctly classifying as S's the stimuli that looked very much like 8's). Mean classification RTs are shown together with classification RTs from Experiment 4b in Figure 20. The relatively slow RTs for the "8" stimuli are due to their infrequent occurrence (16.7% of the trials). Stimulus type continued to be a significant predictor of classification RTs after excluding the "8" trials, $F(4, 40) = 19.38$, $p < .0005$.

4.10.3.3 Relationships between Classification, Search, and the Label Advantage: Individual Differences

Mean individual search RTs and classification RTs were marginally related, $r(9) = .56$, $p = .07$. The participants who were the quickest to classify were also the quickest to find the target. There was no relationship between classification accuracy and overall search RTs.

An exploratory analysis revealed that classification RTs were unrelated to the *label-advantage* measure (described above), but the classification accuracies were[20], though classification errors, particularly for the most typical stimuli were quite low, making these

[20] There was no relationship between classification RT and classification accuracy, thus there is no evidence for a speed-accuracy tradeoff.

analyses subject to ceiling effects. For this reason, I focused on classification accuracy of atypical stimuli. This measure can act as additional proxy measure for category-tuning: individuals with more tightly tuned categories will be more likely to reject atypical members. To determine whether accuracy for the less typical category members correlated with the label-advantage, an average accuracy measure for the three most atypical stimuli was computed and correlated with the magnitude of the label advantage for each subject. The two measures were negatively correlated, $r(9) = -.709$, $p = .015$. Individuals who, after completing the search, had the most trouble classifying the highly atypical S's, were the individuals most benefited by hearing labels[21]. In other words, individuals whose category representations were most strongly centered around the highly typical exemplar (i.e., had the most tightly-tuned category representation) were the individuals benefited by hearing that exemplar labeled.

4.10.4 Discussion

The finding of an association between classification and search RTs is not all that surprising in itself—both are possibly mediated by a general "perceptual speed" factor. What is of more interest is the discovered relationship between classification accuracy for atypical stimuli and the facilitating effect of labels (for the most typical stimuli). Insofar as classification accuracy is a measure of category tuning, the conclusion that can be drawn is that labels most facilitate search for individuals with the most tightly tuned category representations. Supporting this conclusion, an additional measure of category-tuning—the difference in search RTs between the most atypical and most typical trials—revealed that individuals with the largest difference (most tightly-tuned categories) were again the ones most benefiting from hearing labels for the most

[21] Restricting the analysis to the easiest trials only (the condition in which a significant facilitation due to labels was observed), a significant negative correlation was again found between the label advantage score and overall classification accuracy, $r(9) = -.676$, $p = .023$.

typical stimuli. If the two measures of category-tuning are tapping into the same process, they should correlate, and they do, $r(9) = .651$ $p = .030$.

Despite being instructed to think of the highly atypical S stimuli as S's, labels had no effect on these stimuli, even after hundreds of trials, just as being instructed to think of ⊓⊔ as a rotated 2 did not lead to a strong on-line effect of labels (Experiment 1b), but did lead to a sustained facilitation as individuals processed the stimuli as meaningful items rather than novel perceptual symbols—Experiment 2.

If, as predicted by the LFH, the effect of the labels depends on the strength of the association between the stimulus and its category, then establishing a stronger link between the atypical stimuli and category representations should (1) further improve search performance and (2) result in the labels producing an on-line facilitation. It appears that the search task by itself did not establish a sufficiently strong association between the most atypical S stimuli and their categorical representations (despite participants quite accurately classifying them as S's during the training). According to the LFH, this association is necessary for the labels to produce an on-line facilitatory effect. This finding suggest a natural follow-up experiment. If the effect of labels depends critically on a pre-existing association between an item and its category, then exposing individuals to this association by using a category-training session may enable the label to affect even the more atypical stimuli. Insofar as category-training may strengthen the category representation, it should also improve overall search performance[22].

---

[22] One caveat is that strengthening the category representations may lead to a floor effect for the easiest search-trials, eliminating the facilitation due to labels that was observed in this experiment.

4.11 Experiment 4b: Effects of category-training and exploration of individual differences

Experiment 4a led to the following three predictions: (1) providing participants with classification experience should strengthen the degree to which category representations are deployed and thus reduce overall search times. This prediction is not as trivial as it may appear. In Experiment 4a after 400 trials of searching for an identical target—the number 8—there was no hint of a reduction in RTs, $F(19,209) = 1.02$, $p = .44$ (Figure 22), even (especially) for the most difficult search trials, (F's<1). Showing that a brief category training session reduces search times while hundreds of search trials do not would be a noteworthy result. (2) Establishing an association between the atypical S's and the S category during classification training should enable the verbal labels to facilitate search even on the atypical trials. By having participants form an association between these atypical stimuli and the category response (S), hearing "ignore S" may now activate a category representation that includes the atypical S's, thus reducing search RTs for the labeled trials. (3) Insofar as category training conducted immediately before search increases the involvement of category representations, there should be stronger correlations between classification performance, search performance, and on-line influence of labels.

4.11.1 Participants

Twenty-two Carnegie Mellon University undergraduates volunteered for the experiment in exchange for course credit.

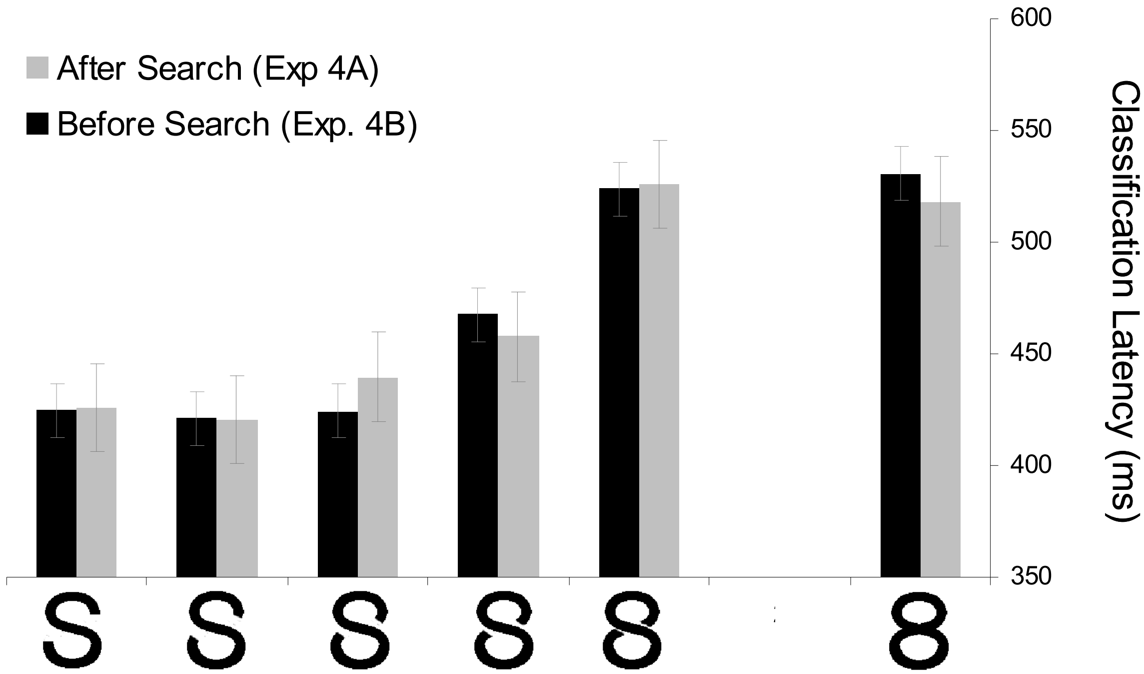

Figure 20: Classification RTs in Experiments 4A (classification after search) and 4B (classification before search). The difference in CIs is due a difference in sample-size between the two experiments.

4.11.2 Stimuli and Procedure

Experiment 4b was identical to 4a except that participants now completed the classification task before the search task. The classification task was meant not only to provide extended exposure to the stimuli prior to beginning the search trials, but also act as a brief classification training session, reflecting the assumption of a close relationship between placing stimuli into meaningful categories, and efficient visual processing. To encourage such training, accuracy feedback was included in the form of a buzz for incorrect responses, and a bleep for correct responses. As in the first set of instructions in Experiment 4a, participants were explicitly instructed that only the complete figure 8's qualify as 8's, and the rest should be classified as S's.

4.11.3 Results

4.11.3.1 Classification/Category Training

An ANOVA with stimulus type as a within-subject factor revealed a significant effect of RTs, $F(5, 105) = 51.22$, $p < .0005$ and accuracy, $F(5, 105) = 24.17$, $p < .0005$. As evident in Figure 20, there were no reliable differences in classification RTs between experiments 4a and 4b. Overall means were equivalent (463ms) and overall accuracy was 96% for both experiments. The correlation between classification RTs in the two experiments for the 5 N-T types was .979. It appears therefore that 400 trials of search made little difference in participants' ability to classify the individual stimuli.

4.11.3.2 Search

The analyses were conducted in a manner identical to Experiment 4a. Non-target (N-T) type and labeling condition were entered into a repeated-measures ANOVA, revealing a highly significant effect of stimulus-type, $F(4, 84) =140.26$, $p < .0005$. Labels marginally facilitated overall search performance, $F(1, 21) = 2.92$, $p = .10$. The label-condition × N-T interaction was highly significant, $F(4, 84) = 3.64$, $p =.009$. Planned comparisons of the effects of labels on RTs for each stimulus type revealed that labels significantly affected search only for the most difficult N-T type, $F(1,21) =7.63$, $p = .012$, resulting in a mean facilitation due to labels of 145 ms. Labels had no effect on target-absent RTs, F<1. The pattern of search RTs is shown in Figure 21.

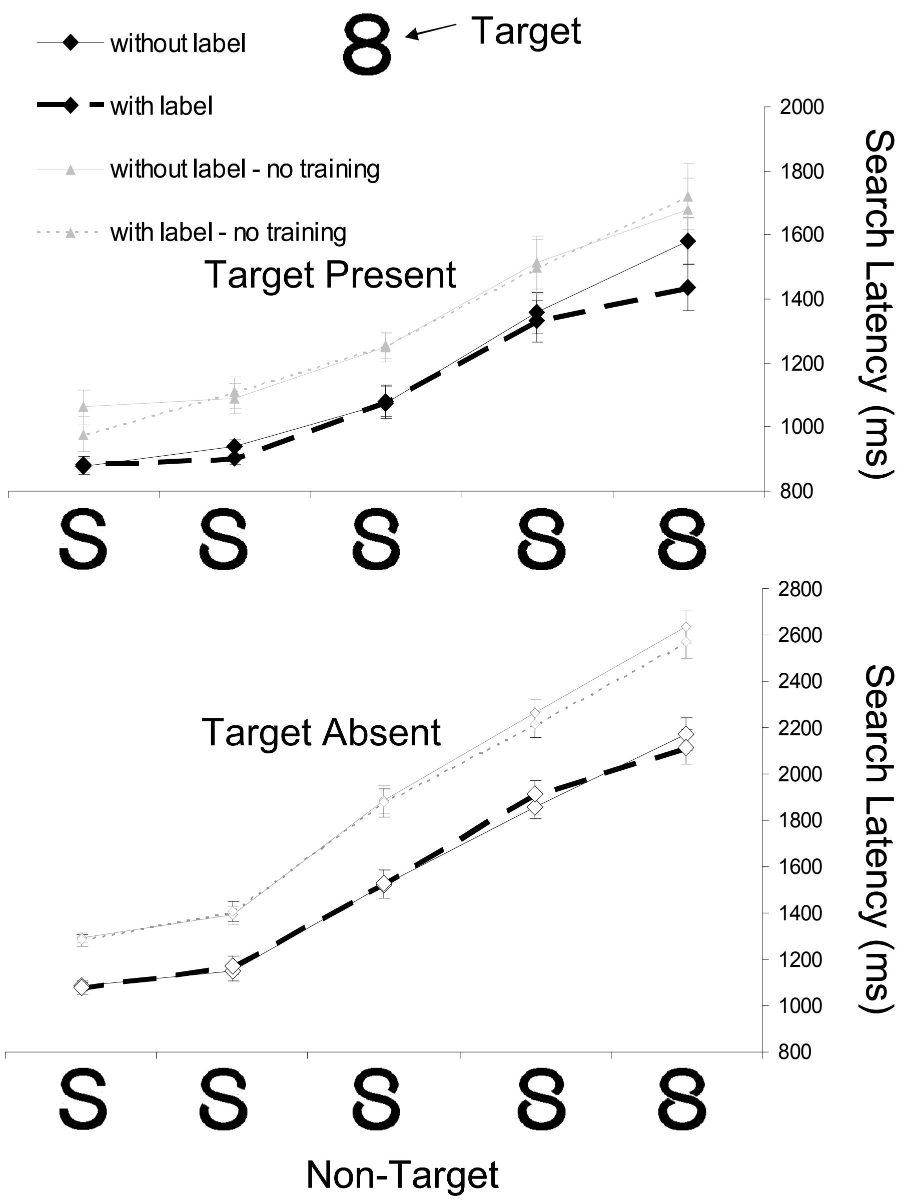

Figure 21: Search performance in Experiment 4b (black lines) and Experiment 4a (gray lines) for comparison. Error bars are within-subject 95% CIs.

Errors increased from 9% on the easiest search trials to 29% on the hardest search trials, $F(4, 84) = 41.56$, $p < .0005$. There were no main effects of the label condition, F<1, but there was a marginal interaction of label condition with N-T type, $F(4, 84) = 2.40$, $p = .06$. For the two most difficult trial-types, labels increase accuracy by 7%, though this advantage did not reach significance, $F(1, 21) = 3.55$, $p = .074$. Interestingly, there was an interaction between the label-accuracy advantage for the more difficult search trials and experiment block (beginning vs. end). In the beginning of the experiment (first 5 blocks), there was no accuracy difference between the

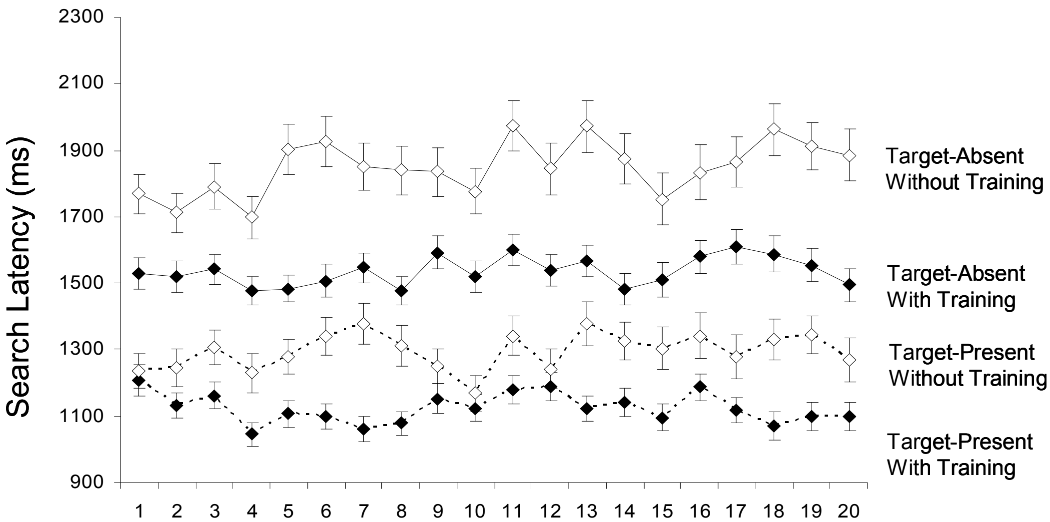

Figure 22: A comparison of search times for Experiments 4a (no training) and 4b (with training). Notice the lack of practice effects within conditions, compared to a dramatic decrease in RTs following the brief category-training session. Error Bars are within-subject 95% CIs.

label and *no-label* trials (about 71%). But by the end (last 5 blocks), accuracy on the *label* trials rose to 81%, while performance on the *no-label* trials fell to 68%, $F(1, 21) = 4.76$, $p < .05$.

4.11.3.3 Effects of Training: A Comparison of Experiments 4a and 4b

As shown in Figure 22, search RTs were consistently lower for participants who first completed the classification session. To compare overall search RTs and accuracy between Experiments 4a and 4b, N-T trial type and block-bin (1=Blocks 1-5; 2=Blocks 6-10; 3=Blocks 11-15; 4=Blocks 16=20) were entered into a mixed-effects ANOVA as within-subject factors and experiment as a between-subject factor.

There was an main effect for overall search times, $F(1, 32) = 8.21$, $p = .007$, as well as a stimulus-type × experiment interaction, $F(4, 32) = 2.98$, $p = .022$. There was no main effect of block-bin, F<1, but interestingly in both experiments, RTs for the most difficult search trials *increased* over the course of the search task by about 200 ms, $F(3, 32) = 3.49$, $p = .019$. An analysis including target-presence as an additional within-subject factor revealed an additional target-presence × experiment interaction, $F(1, 32) = 18.21$, $p < .0005$. Classification training resulted in a greater decrease in RTs on target-absent trials than target-present trials.

There were no differences in accuracy between the two experiments: main effect of RTs, $F<1$; The N-T type × Experiment interaction was likewise not significant, $F(4, 32)<1.93$, $p = .10$, but there was a significant block-bin × experiment interaction, $F(3, 32) = 3.09$, $p = .027$—participants in Experiment 4b started out having slightly lower accuracy (by 3%) than participants in Experiment 4a.

These results can be summarized in the following way. Following approximately 5 minutes of training classifying the target as an 8 and the non-targets as S's resulted in RTs significantly lower than those observed following 35 minutes and 400 search trials. This result is also striking given that RTs did not decrease during the search task. Experience with explicit classification, facilitated discriminating the target from the non-targets, whereas repeated experience with this exact discrimination task for hundreds of trials did not.

4.11.3.4 Relationships between Classification, Label-Effects, and Search Performance: Individual Differences

To investigate the relationship between dependent variables in the classification and search phases of Experiment 4b, dependent variables for the easiest and most difficult search trials were correlated. The theoretically relevant significant and marginally significant correlation values are shown in Figure 23. Care was taken to ensure the correlation coefficients were not inflated by outliers (as defined by having a standardized residual of over 3 standard deviations and/or unusually high leverage). The values in the left side of the figure were computed from the most difficult trials (the non-targets that look very much like S's). The values on the right side were

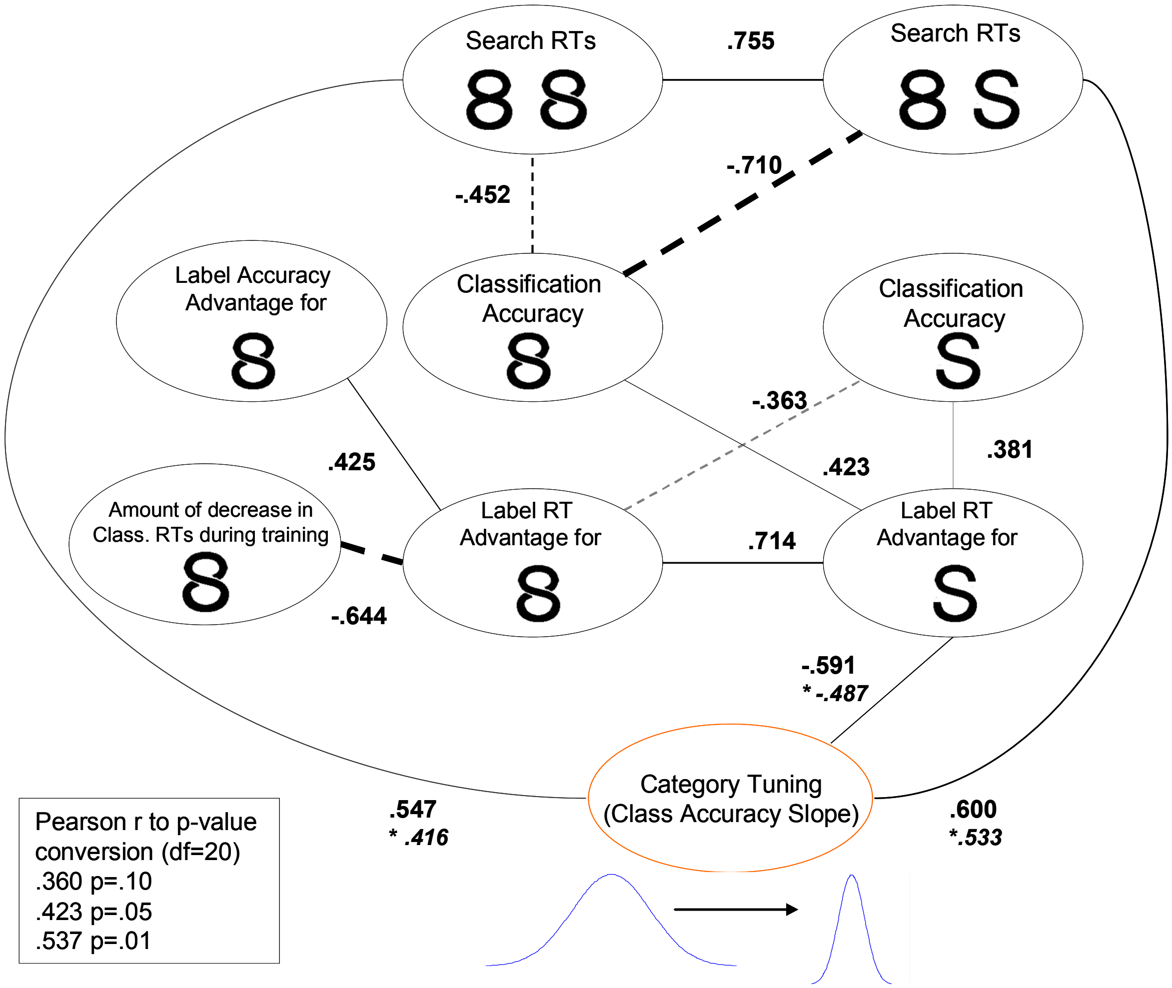

Figure 23: The pattern of correlations between classification, search, and labeling variables in Experiment 4b. Dashed lines indicate negative correlations. Solid lines indicate positive correlations. Line width corresponds roughly to magnitude of the correlation.

*** - partial correlation controlling for average classification accuracy.**

computed from the easiest search trials (the typical S non-targets). The computation of the "Category Tuning" proxy measure is described below.

A primary goal of this experiment was to determine how individual differences in category learning and category representations translate to differences in effects of labels on perceptual processing. To this effect, a new measure of category tuning was obtained by computing the slope of classification accuracy as a function of stimulus typicality. An example of two participants—one with a more broadly tuned category representation than the other—is depicted in Figure 24. A greater slope indicates a more narrowly tuned (and likely overall weaker)

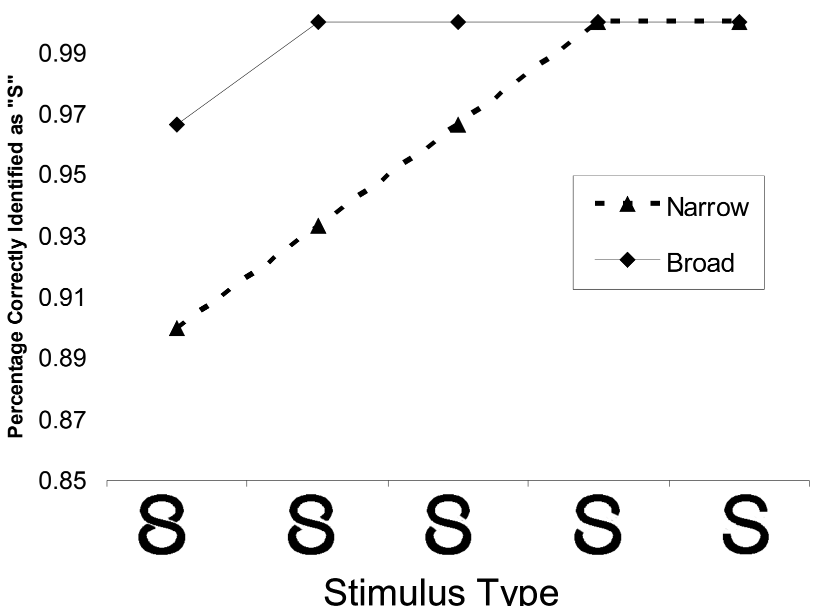

Figure 24: Sample classification accuracy curves for two participants in Experiment 4b. The solid line depicts a more broadly tuned category curve (smaller slope) than the dashed line.

category. One potential problem with this measure is that the slope is confounded with overall accuracy. Thus, individuals with greater slopes tend to have lower overall classification accuracy. To control for this confound, partial correlations were computed that partialed out the overall accuracy, as shown in Figure 23.

The following list outlines the main conclusions drawn from correlation patterns shown in Figure 23.

1. Participants who were benefited by labels on the typical trials also benefited on the atypical trials.
2. Benefits in RT and accuracy due to labels were correlated: participants for whom the labels decreased RTs were the same participants for whom labels increased accuracy.
3. Classification accuracy for the most atypical non-target correlated negatively with search RTs. The higher the classification accuracy (hence broader category tuning), the shorter the RTs.

4. Conversely, more narrowly tuned categories correlated with longer search RTs and a smaller label advantage for the most typical stimuli. These correlations persisted even when overall classification accuracy was partialed out.
5. Greater classification accuracy for the most typical non-target (S) marginally predicted a *greater* label-advantage on the easiest trials, but a *smaller* label-advantage on the atypical trials.
6. The amount of decrease in classification RTs for the most atypical stimuli (RTs for the first 5 blocks – RTs for the last 5 blocks), correlated negatively with the label advantage for the most atypical stimuli. That is, participants who most improved in classifying the atypical stimuli were least benefited by the labels, probably because the category training that immediately preceded search resulted in efficient deployment of category representations even on *no-label* trials.

4.11.4 Discussion

An important motivation for the current experiment was to better understand the mechanisms of the label-facilitation effect through an examination of the factors that predicted this effect for stimuli of different levels of "category goodness." As shown in Experiment 4a, search on the most difficult trials did not seem to be guided by category representations—it was relatively slow and was not affected by category labels.

Experiment 4b demonstrated that a brief classification session significantly facilitated subsequent visual processing of non-targets (as measured by a reduction in search times). This decrease in RTs is striking given that 400 search trials by themselves do nothing to reduce RTs. Completing a brief classification session also resulted in labels reducing search times for the most atypical stimuli. It seems that explicitly establishing an association between atypical category members (stimuli that look more like 8's than S's) and the category / label "S," allowed for the category label to modulate search performance. A measure of category tuning further predicted individual search RTs: individuals with more broadly-tuned, and thus arguably stronger overall category representations had shorter search RTs again illustrating that deploying category representations facilitates search.

According to the LFH, labels modulate visual representations only if an association already exists between the visual stimulus and the label. Combined, Experiments 4a and 4b show that explicit categorization, but not mere exposure to the stimuli, is enough for labels to start to facilitate perceptual processing for atypical category exemplars.

There are several puzzles in the present results. If category training strengthened the category, why was there no longer an effect of labels for the most typical trials? A likely answer is that category training reduced search times on the most typical trials to a floor level beyond which search could no longer be facilitated by category labels. A second puzzle is that proxy measures used for establishing category tuning in Experiment 4a showed that individuals with the most tightly tuned categories benefited *most* from hearing labels. In the present study, individuals with more tightly tuned categories benefited *least* from labels.

One way to resolve this inconsistently is to consider what "category tuning" really is. A more broadly tuned category accepts more atypical items. Thus it would be predicted that insofar as search is guided by category information, individuals with broader categories would perform better on the search task when it includes atypical stimuli. The direction of the correlations between category tuning and search RTs confirms this intuition: more narrowly-tuned categories were correlated with longer RTs, even when controlling for overall classification accuracy.

Insofar as category training with atypical exemplars broadens the tuning of the category, it is predicted to lead to overall better performance, which is evident in Figure 22. What is the relationship between effects of labels and category tuning? Recall that labels are predicted to facilitate search as a function of the association between the category and the stimulus. In an individual with a tightly-tuned category, labels are expected to facilitate processing of the central (most typical), but not peripheral (least typical) exemplars. This is observed in Experiment 4a. To the extent that the brief training session in Experiment 4b strengthened the association between the typical and atypical “S” non-targets and the category/label “S,” we should now observe a label-advantage for the more atypical stimuli, and we do (Figure 21). However, the label-advantage for the typical stimuli now disappears. The most likely explanation for this disappearance, as mentioned above, is a floor effect—the added category/stimulus association strength is now sufficiently strong even without labels. However, there is a caveat. The negative correlation between the category-tuning measure in shown in Figure 23 and the label-advantage for the most typical stimuli indicates that participants with the most broadly-tuned categories continued to derive some benefit from hearing the labels. Further experiments are clearly needed to understand the complex interactions between effects of labels on visual performance and category tuning—measured both prior and post classification training—which not done in the present experiments. The observed differences between experiments 4a and 4b may well be due to the different times at which the classification data is collected: after the search trials in Experiment 4a, and prior to the search trials in Experiment 4b. Future work will investigate the role of category tuning more directly by using more direct measures of individual differences in category structure (Cohen & Nosofsky, 2000).

4.12 Summary

The main prediction of the LFH in the visual domain is that processing an object label should facilitate the processing of the object in a largely automatic way. This prediction was confirmed. Hearing verbal labels such as “two” and “five” prior to the appearance of a search display produced more efficient visual processing as revealed by shorter overall RTs and, importantly, shallower search slopes. While theories like Guided Search predict that seeing (or hearing) a label can facilitate search through top-down guidance (Wolfe et al., 2004), this facilitation is only predicted when the label provides unique information, for instance, telling the participant that on the upcoming trial, she should look for the red vertical (instead of the green horizontal she was searching for on the previous trial). Thus, the effect is not one of language, but of visual cuing. In fact, in this context visual cues tend to work much better than verbal cues.

The observed facilitation due to labels in the present studies occurred even though participants already knew the target and non-targets they need to process on the upcoming trial. Thus the label was redundant in the sense that it was not providing any unique information about the target or non-targets. Labels improved performance even though the target and non-targets were the same for hundreds of trials. In fact in Experiment 4b, hearing labels resulted in *increasing* accuracy on difficult search trials through the course of the experiment, even as performance declined for the randomly intermixed *no-label* trials. The automatic (and temporally transient) nature of the facilitation induced by the labels is made evident in that the reported facilitation is compared to *intermixed* control trials that omit the auditory label, but are otherwise identical. The finding that labeling the (multiple) *non-targets* is equally or more effective than labeling the (single) target is consistent with the interpretation that hearing the label affects processing throughout the visual field.

The facilitation due to labels is highly specific. It is reduced or disappears altogether with rotation of stimuli (Experiments 1a-1b), visual degradation (Experiment 3a), or for more ambiguous/atypical stimuli (Experiment 4a). According to the LFH, this occurs because the manipulations disrupt the association between the exemplar and the stimulus category/label. Thus, better performance on *label* trials compared to the *no-label* control depends on a strong association between the visual exemplar(s) and the labeled category.

This conclusion is supported by two additional sources of evidence, one correlational, and one causal. First, analyses of individual differences (Experiments 4a, 4b) highlighted the strong associations between individual category representations (using several proxy measures of category tuning) and effects of labels. Individuals with the greatest central tendency (most peaked category representations), tended to benefit most from hearing labels (Experiment 4a). The causal connection between strength of exemplar-category association, effects of labels and visual processing are shown in Experiment 4b in which a very brief category training session aimed to strength the association between atypical exemplars the category/category label, resulted in better performance than could be achieved by practice alone. Strengthening the association between the category and atypical exemplars now also allowed the label to facilitate their processing.

In summary, labels, categories, and visual processing appears to interact on two levels. First, there is a sustained benefit to ascribing meaning to a perceptually novel stimulus (i.e., placing it into a meaningful category) as shown in Experiment 2. Thus, difficulties in processing perceptually novel items arise not only from their novelty, but also from a lack of top-down support from their category. Second, beyond this sustained facilitation, labels produce further transient *on-line* facilitation that is highly sensitive to the typicality of the stimulus.

5 Conclusion

This thesis began with the question "What does language do?" The traditional view is that language functions for communicating our ideas to others, but does not alter the content of the ideas (their representations) in any significant way (Bloom & Keil, 2001; Gleitman & Papafragou, 2005; Li & Gleitman, 2002; Pinker, 1995). This thesis presented evidence from a number of domains that is inconsistent with this view. Most of the reviewed studies—from development, aphasia, and cross-linguistic/cultural experiments—are limited in their ability to establish causal connections of language on mental representations. For instance, finding that Indonesian speakers are less likely to remember whether a picture they saw depicted an action in the past or future (Boroditsky et al., 2002) points to a link between morphological tense markers—required in English, but absent in Indonesian—and memory. But such a finding by itself, tells us little about the mechanism by which language might affect memory. Perhaps, continued practice with having to attend to tense—probably for future communication—has simply resulted in automatic encoding of tense (Levinson, 2003) or a special "thinking for speaking" mode (Slobin, 1996). While this does not diminish the importance of such findings—whatever the mechanism, Indonesian speakers tend to have poorer memory for the time of events than English speakers—it does not tell us whether language plays an active role in building or modulating mental representations.

Studies showing that verbal interference partially removes these types of linguistic effects (and more dramatically, affects perceptual judgments, see Section 2.2) go further in illuminating the mechanisms possibly at work. Together with findings showing that there appears to be a link between language loss and performance on certain cognitive skills not requiring communication (categorization most notably, see Section 2.5), this evidence suggests that language, and verbal labels in particular have on-line effects on cognition in tasks previously considered to be independent of language (e.g., sorting colors).[23]

What is missing from the now-numerous studies detailing effects of language on various cognitive skills, is a description of a mechanism by which language may come to have these effects. This thesis attempted to provide such a mechanism in the form of the Language Feedback Hypothesis (Section 3). Rather than limiting investigations of effects of labels to the kinds of high-level conceptual representations tapped into by categorization tasks (e.g., the alien-categorization experiment described in Section 2.4), I have attempted to go deeper, predicting that effects of language on categorization might have a more fundamental basis in perceptual processing.

I have argued that insofar as verbal labels become associated with clusters of correlated feature, the activation of a label activates the corresponding visual features through a process of neural feedback. As argued in Section 3.3, feedback from higher-level to lower-level cortical regions is not only ubiquitous in the visual system, but are possibly *required* for visual awareness (3.3.3). Conceptualizing effects of language on visual processing as a feedback process, explains why hearing a label for a familiar and overlearned item like "2" can affect its visual processing.

Critically, the notion of language affecting representations online is not limited to visual processing, and can be applied to categorization, memory, *and any other process that takes visual information as its input*. Indeed, this is the precise mechanism proposed for the finding that calling common objects by their names reduces within-category recognition memory (2.6.2).

Much, much more work is needed for understanding the implications of these findings. Languages differ in what object and event categories are labeled, and how the labels map onto

---

[23] My own verbal interference study mimicked one pattern of impairments observed in aphasia, further arguing that impeding language functions produces certain types of categorization impairments (2.6.1).

the various objects and events. Does experience with different patterns of labeling change visual processing? Does language enable us to make the novel meaningful? Recall that referring to a  as a rotated 2 facilitate visual processing—more than was observed with simple practice. To what degree does the *linguistic* ability to combine dimensional adjectives like “gray small vertical” enable us to deploy attention to such visual conjunctions? Indeed, to what degree do dimensional terms like “size” and “color” *enable* us to selectively attend to these attentions? Would speaking a language that does not contain terms like “green” and “square” but contains a term that means a “green square” lead its speakers to be particularly efficient at attending to and processing green squares, but at the expense of selectively attending to green-ness and square-ness in other contexts?

Human cognition is characterized by its flexibility (Barsalou, 1983; Deak, 2003). We can, at will, attend to, and group, objects by color, or size, or by whether they are appropriate things to bring on a picnic. We have little trouble ignoring physical differences, grouping together “A” and “a” or Chihuahuas and German Shepherds, but when necessary, have no trouble telling them apart. This intense flexibility may be in part produced by language. The studies presented in this thesis open the door to understanding the mechanisms by which language contributes to the unique aspects of human cognition.

6 Appendix

6.1 *The Conceptual Grouping Effect: Categories Matter (and named categories matter more)*

People interpret what they see—quickly and automatically recognizing familiar objects as members of categories (Grill-Spector & Kanwisher, 2005). To what degree is visual processing itself shaped by conceptual knowledge? The classic separation between perceptual and conceptual systems has been challenged by mounting evidence for a much more interactive view (for review see Goldstone & Barsalou, 1998). Evidence from single-cell recordings has further blurred the line between the bottom-up processes of "pure" perception and top-down feedback that is potentially open to conceptual influences (Hupe et al., 2001; Lamme, Super, & Spekreijse, 1998; Lee & Nguyen, 2001). The remarkable speed at which object categorization occurs (Fabre-Thorpe, Delorme, Marlot, & Thorpe, 2001) further suggests that basic perceptual processes such as attentional selection and grouping may be penetrable by conceptual knowledge. While much is known about the effects of category-learning on perceptual organization, for example, the improved ability to discriminate stimuli following category training (e.g., Goldstone, 1994), considerably less is known about how object categories influence perceptual processes on-line. The present experiments use the paradigm of visual search to study how what we know affects what we see.

Theories of visual processing have often overlooked the possible contributions of conceptual categories (Wolfe & Horowitz, 2004 for discussion). In particular, the idea that conceptual categories affect performance in the domain of visual search has fallen into disfavor following failures to replicate Jonides and Gleitman's (1972) oh-zero effect (e.g., Duncan, 1983) and findings arguing that category effects hinge on perceptual rather than conceptual factors (Krueger, 1984; Levin et al., 2001). At the same time, it is clear that visual search performance cannot always be reduced to low-level visual factors. It is strongly affected by familiarity (Frith, 1974; Malinowski & Hubner, 2001; Wang et al., 1994) and controlling for physical differences, is sensitive to the categorical relationship between targets (T's) and non-targets (N-T's) such as "blue vs. green" (Daoutis et al., 2006) and "steep vs. non-steep" (Wolfe et al., 1992). The origin, mechanisms, and specificity of these effects remain largely unknown. For instance, it is unclear to what degree representational differences between color categories and tilt categories are the product of experience and learning versus physiological constraints, and to what degree they are due to long-term representational change versus on-line representational reorganization.

The first aim of this work was to test for the presence of category effects in visual search while (1) controlling for all physical factors and (2) using familiar yet clearly learned stimuli. This was achieved by varying the *conceptual heterogeneity* of letter non-targets. It is known that physical N-T heterogeneity correlates positively with search times—searching for a T among L's is harder if L's are presented in varying orientations due to grouping of perceptually similar N-T's (Duncan & Humphreys, 1989). Experiment 1 tests for the presence of *conceptual grouping* by investigating whether N-T heterogeneity similarly slows search.

An effect of conceptual categories on visual processing can be attributed to two sources. First, items within a conceptual category may have become represented as more similar due to extensive practice with categorizing together these stimuli (e.g., Goldstone, 1994; Harnad, 1987). In this way, conceptual homogeneity may have turned into perceptual homogeneity. Alternatively, conceptual grouping may arise dynamically, through top-down modulation of visual representations by category-level representations. Experiments 2 and 3 examined these possibilities.

6.1.1 Experiment 1

The main goal of this experiment was to test for the existence of *conceptual grouping* effects by varying conceptual non-target heterogeneity (between-category versus within-category) among perceptually equidistant non-targets. A secondary goal was to examine the specificity of this effect by manipulating the distinctiveness of the target and the familiarity of the non-targets.

Methods

Participants

Twenty-one Carnegie Mellon University undergraduates (aged 18-22) volunteered for the experiment in exchange for course credit or $7. Two participants were eliminated for having accuracy below 80%.

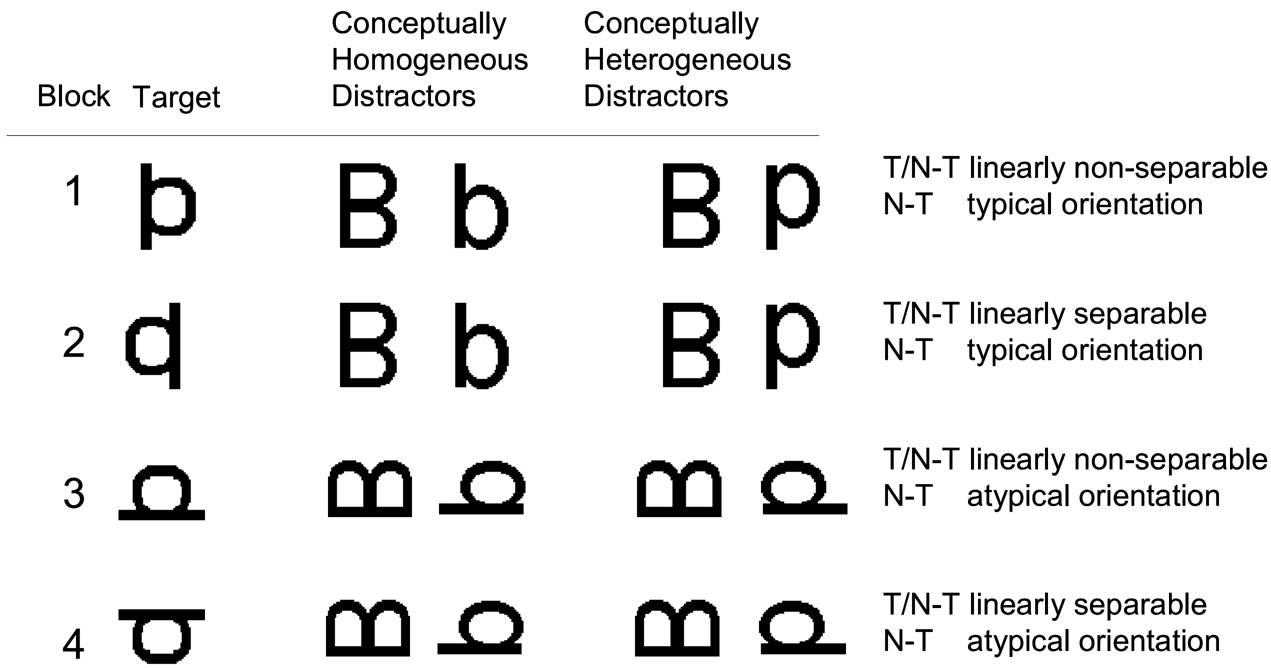

Figure 25: The assignment of targets (T) and non-targets (N-T) to blocks in Experiment 1. Linear separability refers to a difference in gross orientation between T and N-Ts. All conditions are within-subject with participants completing the blocks in counterbalanced order. Within each block, the two trial-types (homogeneous and heterogeneous) were intermixed.

Stimuli and Procedure

Participants completed four blocks in counterbalanced order, searching for a non-letter target among conceptually heterogeneous non-targets (B and p) or conceptually homogeneous non-targets (B and b)[24]. The full assignment of targets and non-targets to blocks is shown in Figure 25. Each character subtended .7° × .8°. of visual angle. The characters were white, displayed on a black background and arranged along the circumference of an imaginary circle (7° diameter) around a fixation cross (.5° diameter). The placement of the target and non-targets was random with the stipulation that the same number of items were present on the left and right sides of the display.

Each block consisted of 12 practice trials followed by 9 repetitions of 12 trials (target-present versus target-absent × 3 display sizes—4, 6, 10 × within-category versus between-category N-

---

[1] A capital “B” is often rendered with its lower-loop larger than the upper-loop. The B used in these studies was modified to ensure both loops were identical.

T's). Trial order was random with the target present on exactly half the trials. Participants gave 2-alternative target *present* / *absent* responses using a gamepad controller. Participants were instructed to respond as quickly as possible without compromising accuracy. If accuracy dipped below 92% for 24 trials, participants were prompted to try to be more accurate. Response mapping of right/left hands to present/absent responses was counterbalanced between participants. The inter-trial interval was 750 ms. Feedback in the form of a buzzing sound was provided for incorrect responses.

Results and Discussion

Analyses will focus on target-present trials. Incorrect responses and responses shorter than 150 ms or longer than 3,500 ms were excluded. Search in blocks in which the T and N-T faced in different directions (were linearly separable) was much faster, $F(1,18) = 114.51$, $p < .0005$ and more efficient, $F(2, 36) = 37.85$, $p < .0005$ than when T and N-T faced in the same direction (linearly nonseparable). Search through upright letters was faster, $F(1,18) = 19.18$, $p < .0005$ and more efficient, $F(2, 36) = 3.82$, $p = .03$, than search through rotated letters (Figure 26). The finding of faster search through familiar than unfamiliar non-targets is hardly new, but what is significant is that the present finding cannot be attributed to a difference in novelty between target and non-targets (Malinowski & Hubner, 2001; Treisman & Gelade, 1980; Wang et al., 1994)—since the target was always novel—supporting the interpretation that such effects have more to do with greater processing efficiency of familiar stimuli than differences in familiarity between Ts and N-T's (Rauschenberger & Yantis, 2006; Richards & Reicher, 1978). Interestingly, the effect of linear-separability was itself mediated by familiarity: linearly nonseparable search was particularly difficult when the N-T's were in a less-familiar orientation;

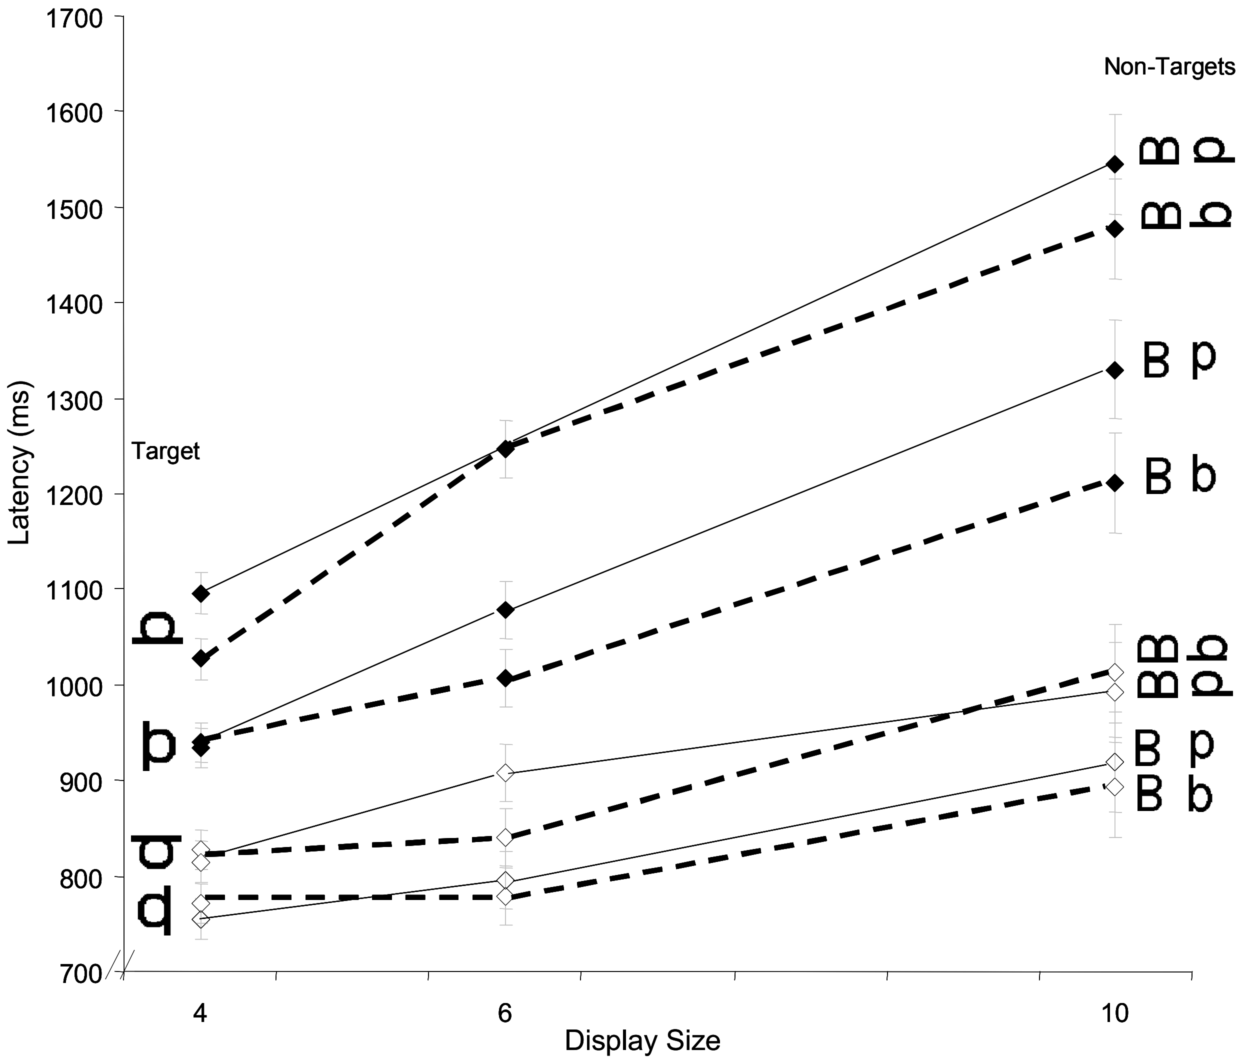

Figure 26: Search times by block and trial-type, as a function of display size in Experiment 1. Only target-present trials are shown. Solid symbols indicate linearly non-separable blocks. Bars show within-subject 95% confidence intervals based on condition interactions (Loftus & Masson, 1994).

linear nonseparability was less detrimental for letters in their canonical orientation, $F(1,18) = 9.98$, $p = .005$. These differences in RTs were paralleled by error data which showed no evidence of a speed-accuracy tradeoff in any of the blocks.

The subsequent analyses focus on the conceptual relationship between non-targets. A comparison of search through within- and between-category non-targets revealed significantly faster RTs for same-category non-targets, $F(1,18) = 6.27$, $p = .02$. There was no overall difference in error rates, $F < 1$. Planned comparisons revealed that the conceptual relationship between the

N-T affected performance most while searching for þ among upright letters. In this block, search through within-category non-targets was both faster, $F(1,18) = 6.22$, $p = .02$, and more efficient, $F(2,36) = 5.21$, $p = .01$, reducing search slopes from 65 ms/item to 47 ms/item (Figure 26). There were no overall differences in error rates, $F<1$, though, accuracy (in Block 1 only) was significantly higher for the smallest display size of the within-category (B-b) condition, $F(1,18) = 4.18$, $p<.05$. The effect of conceptual homogeneity did not reach significance for the other blocks.[25] Together, these results provide evidence for a conceptual grouping effect for familiar stimuli, particularly when discriminating between T and N-T's is difficult.

6.1.2 Experiment 2

There are several possible explanations for the conceptual grouping effect observed in Experiment 1. Given years of experience categorizing B's and b's as members of the same category, the two letters may have come to look increasingly similar to each other—a type of categorical perception (or perceptual warping) effect (Goldstone, 1994; Kuhl, 1994). Alternatively, the conceptual grouping effect may emerge on-line during the search task possibly due to top-down effects of category-level information. The purpose of Experiment 2 was to evaluate the first alternative by using a speeded same-different judgment task. If B is more similar to b than p, then one should observe slower RTs in judging of physical difference of B-b pairs compared to B-p pairs. A second goal of Experiment 2 was to ensure that the difference between within- and between-category non-targets in Experiment 1 was not due to unforeseen perceptual confounds, such as the target þ being represented as more similar to a b than a p (in which case the observed effect could be attributed directly to differences in T/N-T similarity).

[25] However, as might be expected, search on target-absent trials of Block 3 showed significantly faster search through conceptually homogeneous N-T's, $F(1,18) = 8.49$, $p = .01$.

Methods

Participants

Fourteen Carnegie Mellon undergraduates participated for course credit.

Stimuli and Procedure

The stimuli were identical to Experiment 1. A speeded same/different judgment task was used. Participants were presented with T/N-T combinations used in Experiment 1. The pairs were randomly placed to the left or right of fixation and participants were instructed to respond "same" only if the stimuli were physically identical. Participants completed 8 practice trials followed by 120 same-different judgments containing 15 repetitions of the comparisons shown in Figure 3. Trial-types were intermixed. Feedback in the form of a buzzing sound was provided for incorrect responses.

Results and Discussion

Mean RTs and statistical comparisons are shown in Figure 27. There were no differences in RTs for responses to B-b and B-p pairs, arguing against the idea that perceptual-warping is the source of the conceptual grouping effect. There were also no differences in RTs between þ-p and þ-b responses, confirming that the non-letter target þ was perceptually equidistant from both "b" and "p" non-targets, thus ruling out any spurious perceptual confounds in Experiment 1. There was a very substantial effect of familiarity. The RTs for responding "same" to þ-þ were more than 200 ms longer than for making "same" responses to familiar letters. Of interest also is that B-B comparison was faster than b-b / p-p comparison, further adding to the literature on effects of symmetry in visual perception (Richards, 1978).

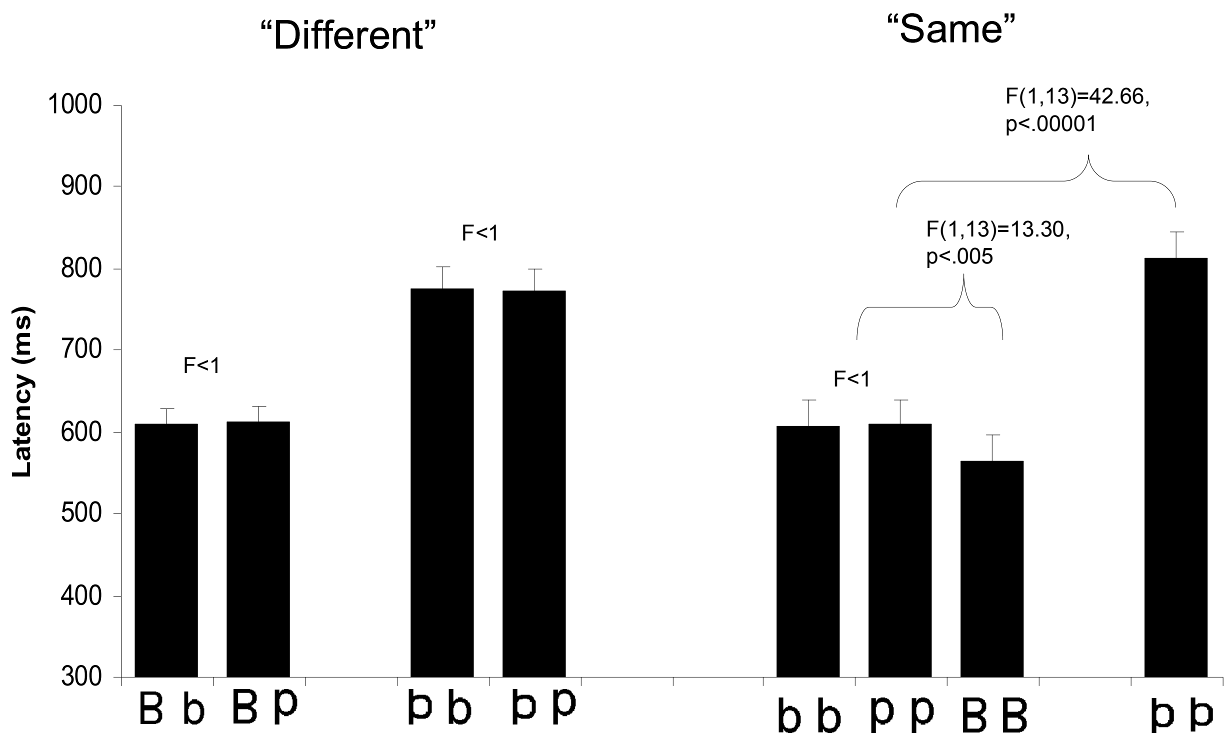

Figure 27: Reaction-times for same/different judgments in Experiment 2. Participants judged the pairs for physical identity. Each bar shows the mean RT for the pair indicated. Error bars show within-subject 95% CIs.

The failure to find significant differences between B-p and B-b judgments does not support a perceptual-learning account in which the conceptual-grouping effect can be explained through long-term changes to representations of stimuli in the same conceptual category. The possibility that conceptual effects in visual search arise on-line through top-down modulation of visual representations with conceptual knowledge was tested in Experiment 3.

6.1.3 Experiment 3

Experiment 3 examined the impact of verbal category labels on visual search. Consider that the visual stimulus "B" is not just a member of a familiar category (one that can be instantiated using a variety of perceptual forms: b, *b*, B), it is a member of a named category. Over time, category labels (i.e., "bee") become strongly associated with features that are most diagnostic (or typical) of the named category. If conceptual categories affect visual processing on-line, then hearing a category name prior to the appearance of a search display may further modulate the degree to which visual representations are shaped by conceptual categories. The purpose of Experiment 3 was to determine whether cuing a target (B) with a verbal label ("bee") facilitates search over and above simply knowing what the target is. A facilitation effect would support the hypothesis that conceptual category effects results from on-line top-down modulation. Insofar as verbal labels are associated with category exemplars, hearing a label may allow conceptual categories to further penetrate visual processing. This effect should be sensitive to task requirements. Hearing a label ("find the B") may facilitate search when T and N-T's are in different conceptual categories (B-p), but hinder search when they are in the same category (B-b).

Methods

Participants

Twenty-eight Carnegie Mellon undergraduates participated for course credit. Data for one participant was missing for Block 2 due to experimenter error.

Stimuli and Procedure

Participants completed two blocks in counterbalanced order: In Block 1, they searched for a "B" among "b" N-T's (within-category trials) or "p" N-T's (between-category trials). Prior to each search trial, the target ("B") was either verbally cued (*label* condition) or not. Target identity was always known, so the label did not add any additional information. To assess the specificity of any label-effects, an additional block of trials maintained all the low-level properties of the original "B" target, but mirror-reversed it ("ᗺ"), thus arguably disrupting or weakening the association between it and the verbal label. Each trial started with a fixation cross

(500 ms) followed by an auditory prompt ("find the B" or "find the target") (1000 ms). The search display appeared 600 ms after the offset of the verbal prompt. For each block, participants completed 12 practice trials followed by 10 repetitions of 24 trials (target present versus target absent × 3 display sizes × labeling condition × trial-type: within-category or between-category). The procedure was otherwise identical to Experiment 1.

Results and Discussion

Search was highly efficient (<5 ms/item) so the analyses collapse across display size. Hearing "find the B" prior to the appearance of the search display facilitated performance on between-category trials only as revealed by a significant labeling × trial-type interaction, $F(1, 27) = 4.38$, $p = .05$. When searching for a B among p's, labels significantly reduced search times, $t(27) = 3.44$, $p = .002$ (Figure 28-left)**.** This label-facilitation effect was highly specific, showing a very different pattern of results when the target was mirror-reflected (Figure 4-right). Now, labels facilitated search only when the *non-targets* were lowercase b's, $t(26) = 2.22$, $p = .04$. The labeling × trial-type interaction was again significant, $F(1,26) = 8.07$, $p = .01$. There were no accuracy labeling × trial-type interactions in either block, F's<1. There were no overall differences in RTs or errors between blocks, F's<1.

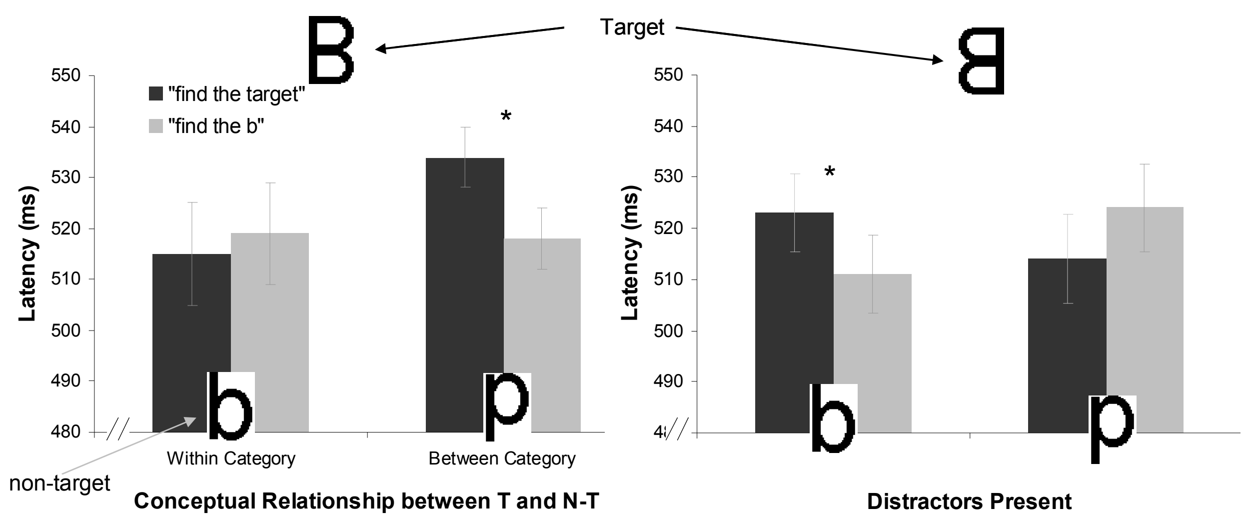

Figure 28: The effect of labels on search times for a familiar target B (left), and for a mirrored B (right) in Experiment 3. Labels facilitate search when the task requires discriminating a B from a non-B. Bars show within-subject 95% confidence intervals. Asterisks signify significant differences between means at *p*<.05.

A parsimonious explanation of this pattern of results is that actually hearing the category label facilitates performance whenever the task requires discriminating a B (lowercase or uppercase) from a non-B. In contrast, no facilitation due to the label is observed in trials requiring discrimination within a conceptual category—B's and b's (indeed, there is a slight though not significant cost in both RTs and accuracy), or when no b's are present in the display (searching for a ꓭ among p's)—Figure 28-right**.** This last result highlights the importance of a pre-existing association between the label and the visual stimulus for obtaining the label-facilitation effect. Unlike Experiment 1, in which conceptual effects were most clearly observed only in a difficult search (>40 ms/item), here, labels penetrated visual processing of a highly efficient search.

A curious aspect of the present results is that in the *no-label* condition, within-category search was actually easier than between-category search—a finding that seems at odds with work showing a superiority for between-category search for colors (Daoutis et al., 2006; Gilbert et al., 2006). A possible explanation is that the activation of the category representation by the category labels acted as a sustained prime, leading to overall faster processing of b's overall (even while making it more difficult to discriminate between B's and b's). A replication of Experiment 3 without labels at all (*N*=23) supported this interpretation. While search was still highly efficient (<5ms/item), overall RTs were significantly slower than in Experiment 3 (*M*=584 ms *SD*=104 ms versus *M*=521 ms, *SD*=58 ms), $t(32) = 2.57$, $p = .02$. Consistent with the idea that faster within-category search in Experiment 3 was due to the presence of labels, there was now no difference between within- and between-category search, $F<1$. It appears then, that hearing the category name not only facilitated search on between-category *label* trials, but also had a more sustained impact of making the processing of b's more efficient, perhaps by making their representations more informationally redundant (Garner & Clement, 1963; Rauschenberger & Yantis, 2006).

General Discussion

Conceptual categories affected visual search performance as revealed by faster search times through within-category (conceptual homogeneous) N-T's compared to between-category (conceptual heterogeneous) N-T's. This *conceptual grouping effect* seems to arise on-line, perhaps through top-down feedback of category-level representations onto lower-level visual representations (Lupyan, 2007), rather than through pre-existing differences in similarities between stimuli in the same versus different conceptual categories. Supporting this claim, Experiment 2 failed to find differences between responses to within-category B-b pairs and between-category B-p pairs, lending support to the idea that the conceptual grouping effect observed in Experiment 1 emerges on-line during the search task (i.e., is a true grouping effect) rather than a result of prior experience with the letters permanently changing the perceptual space (Kuhl, 1994). In further accord with this interpretation, verbal labels enhanced the degree to which conceptual categories penetrated perceptual processing. In Experiment 3 it was shown that in a mixed-trial design, search times were reduced when a target was labeled compared to trials on which it was not (with target-identity always known). This facilitation due to labels was highly specific to stimuli that had pre-existing associations with the label—B and b, but not p or ᗺ.[26]

What mechanism might be responsible for the finding that simply hearing labels—which contribute no additional knowledge—facilitates between-category search? Most compatible with the present results are theories that stress the fluid interaction between higher- and lower-levels of visual processing such as Hochstein and Ahissar's (2002) Reverse Hierarchy Theory. In accord with this theory, verbal labels may engage higher-levels of visual representations than are engaged in the absence of labels. These more categorical representations facilitate search by dynamically collapsing low-level differences within a category (conceptual grouping), while exaggerating the representational differences between the named category and other stimuli.

[26] The finding that hearing "find the b" also facilitates searching for a ᗺ among b's, suggests that explicitly labeling the non-targets (e.g., "ignore the b's") should also facilitate search. This has been confirmed by Lupyan (2007b), who also found reduced search slopes on label trials in a more difficult task (searching for a 2 among 5's).

7 References

Arnoult, M. D. (1953). Transfer of Predifferentiation Training in Simple and Multiple Shape Discrimination. *Journal of Experimental Psychology, 45,* 401-409.

Balaban, M. T. & Waxman, S. R. (1997). Do words facilitate object categorization in 9-month-old infants? *Journal of Experimental Child Psychology, 64,* 3-26.

Baldwin, D. A. & Markman, E. M. (1989). Establishing Word-Object Relations - A 1St Step. *Child Development, 60,* 381-398.

Barsalou, L. W. (1983). Ad Hoc Categories. *Memory & Cognition, 11,* 211-227.

Bartlett, F. C. (1932). *Remembering*. Cambridge, England: Cambridge University Press.

Battig, W. F. (1956). Transfer from Verbal Pretraining to Motor-Performance As A Function of Motor Task Complexity. *Journal of Experimental Psychology, 51,* 371-378.

Behl-Chadha, G. (1996). Basic-level and superordinate-like categorical representations in early infancy. *Cognition, 60,* 105-141.

Berlin, B. & Kay, P. (1969). *Basic Color Terms: Their Universaility and Evolution*. Berkeley and Los Angeles: University of California.

Bloom, P. & Keil, F. C. (2001). Thinking through language. *Mind & Language, 16,* 351-367.

Booth, A. E. & Waxman, S. (2002). Object names and object functions serve as cues to categories for infants. *Developmental Psychology, 38,* 948-957.

Boroditsky, L. (2003). Linguistic Relativity. In L.Nadel (Ed.), *Encyclopedia of Cognitive Science* (pp. 917-922). London: Macmillan.

Boroditsky, L. (2000). Metaphoric structuring: understanding time through spatial metaphors. *Cognition, 75,* 1-28.

Boroditsky, L., Ham, W., & Ramscar, M. (2002). What is universal about event perception? Comparing English and Indonesian speakers. In *Proceedings of the 24th Annual Meeting of the Cognitive Science Society* Fairfax, VA.

Bowerman, M. (1989). Learning a semantic system: What role do cognitive predispositions play? In M.Rice & R. Schiefelbusch (Eds.), *The Teachability of Language* (Baltimore, MD: Paul H. Brooks.

Bowerman, M. (1996). Learning how to structure space for language: A cross-linguistic perspective. In P.Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and Space* (Cambridge, MA: MIT Press.

Bowerman, M. & Choi, S. (2001). Shaping meanings for language: universal and language-specific in the acquisition of spatial semantic categories. In M.Bowerman & S. C. Levinson (Eds.), *Language acquisition and conceptual development* (pp. 475-511). Cambridge, UK: Cambridge University Press.

Bransford, J. D. & Franks, J. J. (1971). Abstraction of Linguistic Ideas. *Cognitive Psychology, 2,* 331-350.

Brown, R. (1986). *Social Psychology. The Second Edition*. New York: Free Press.

Buccino, G., Riggio, L., Melli, G., Binkofski, F., Gallese, V., & Rizzolatti, G. (2005). Listening to action-related sentences modulates the activity of the motor system: A combined TMS and behavioral study. *Cognitive Brain Research, 24,* 355-363.

Caramazza, A. & Berndt, R. S. (1982). The Semantic Deficit Hypothesis - Perceptual Parsing and Object Classification by Aphasic Patients. *Brain and Language, 15,* 161-189.

Carlson, L. A. & Logan, G. D. (2001). Using spatial terms to select an object. *Memory & Cognition, 29,* 883-892.

Carruthers, P. (2002). The Cognitive Functions of Language. *Behavioral and Brain Sciences, 25,* 657-674.

Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). A Neural Basis for Visual-Search in Inferior Temporal Cortex. *Nature, 363,* 345-347.

Choi, S. (2006). Influence of language-specific input on spatial cognition: Categories of containment. *First Language, 26,* 207-232.

Christiansen, M. H. & Chater, N. (1999). Connectionist natural language processing: The state of the art. *Cognitive Science, 23,* 417-437.

Clark, A. (1998). Magic Words: How Language Augments Human Computation. In P.Carruthers & J. Boucher (Eds.), *Language and Thought: Interdisciplinary themes* (pp. 162-183). Cambridge University Press.

Clark, A. & Karmiloff-Smith, A. (1993). The Cognizer's Innards: A Psychological and Philosophical Perspective on the Development of Thought. *Mind & Language, 8,* 487-519.

Clark, E. V. (1973). What's in a word? On the child's acquisition of semantics in his first language. In T.E.Moore (Ed.), *Cognitive Development and the Acquisition of Language* (pp. 65-110). New York: Academic Press.
Cohen, A. L. & Nosofsky, R. M. (2000). An exemplar-retrieval model of speeded same-different judgments. *Journal of Experimental Psychology-Human Perception and Performance, 26,* 1549-1569.
Cohen, R., Kelter, S., & Woll, G. (1980). Analytical Competence and Language Impairment in Aphasia. *Brain and Language, 10,* 331-347.
Cohen, R., Woll, G., Walter, W., & Ehrenstein, H. (1981). Recognition deficits resulting from focussed attention in aphasia. *Psychological Research, 43,* 391-405.
Colunga, E. & Smith, L. B. (2005). From the lexicon to expectations about kinds: A role for associative learning. *Psychological Review, 112,* 347-382.
Corthout, E., Uttl, B., Walsh, V., Hallett, M., & Cowey, A. (1999). Timing of activity in early visual cortex as revealed by transcranial magnetic stimulation. *Neuroreport, 10,* 2631-2634.
Cromer, R. F. (1974). The development of language and cognition: The cognition hypothesis. In B.Foss (Ed.), *New perspectives in child development.* (Harmondsworth: Penguin.
Daoutis, C. A., Pilling, M., & Davies, I. R. L. (2006). Categorical effects in visual search for colour. *Visual Cognition, in press*.
Davidoff, J. (2001). Language and perceptual categorization. *Trends in Cognitive Sciences, **5**,* 382-387.
Davidoff, J., Davies, I., & Roberson, D. (1999). Colour categories in a stone-age tribe. *Nature, 398,* 203-204.
Davidoff, J. & Roberson, D. (2004). Preserved thematic and impaired taxonomic categorisation: A case study. *Language and Cognitive Processes, 19,* 137-174.
Deak, G. O. (2003). The development of cognitive flexibility and language abilities. In R.Kail (Ed.), *Advances in Child Development and Behavior Vol. 31* (pp. 271-327). San Diego: Academic Press.
Dell, G. S. (1986). A Spreading-Activation Theory of Retrieval in Sentence Production. *Psychological Review, 93,* 283-321.
Dell, G. S. & Oseaghdha, P. G. (1992). Stages of Lexical Access in Language Production. *Cognition, 42,* 287-314.
Dell, G. S., Schwartz, M. F., Martin, N., Saffran, E. M., & Gagnon, D. A. (1997). Lexical access in aphasic and nonaphasic speakers. *Psychological Review, 104,* 801-838.
Derivera, J. (1959). Some Conditions Governing the Use of the Cue-Producing Response As An Explanatory Device. *Journal of Experimental Psychology, 57,* 299-304.
Desimone, R., Albright, T. D., Gross, C. G., & Bruce, C. (1984). Stimulus-Selective Properties of Inferior Temporal Neurons in the Macaque. *Journal of Neuroscience, 4,* 2051-2062.
Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: The psychophysics of reentrant visual processes. *Journal of Experimental Psychology-General, 129,* 481-507.
Dick, A. O. (1971). Processing Time for Naming and Categorizatoin of Letters and Numbers. *Perception and Psychophysics, 9,* 350-352.
Diesendruck, G. (2003). Categories for names or names for categories? The interplay between domain-specific conceptual structure and language. *Language and Cognitive Processes, 18,* 759-787.
Donderi, D., Seal, S., & Covit, L. (1973). Association Learning and Visual Discrimination. *Perception & Psychophysics, 14,* 394-400.
Donderi, D. C. & Kane, E. (1965). Perceptual-Learning Produced by Common Responses to Different Stimuli. *Canadian Journal of Psychology, 19,* 15-30.
Drivonikou, G. V., Kay, P., Regier, T., Ivry, R. B., Gilbert, A. L., Franklin, A. et al. (2007). Further evidence that Whorfian effects are stronger in the right visual field than the left. *Proceedings of the National Academy of Sciences of the United States of America, 104,* 1097-1102.
Druks, J. & Shallice, T. (1996). Modality specific naming impairments: The case of multiple semantics. *Brain and Cognition, **32**,* 245-247.
Druks, J. & Shallice, T. (2000). Selective preservation of naming from description and the "restricted preverbal message". *Brain and Language, 72,* 100-128.
Duncan, J. (1983). Category Effects in Visual-Search - A Failure to Replicate the Oh-Zero Phenomenon. *Perception & Psychophysics, 34,* 221-232.
Duncan, J. & Humphreys, G. W. (1989). Visual-Search and Stimulus Similarity. *Psychological Review, 96,* 433-458.
Fabre-Thorpe, M., Delorme, A., Marlot, C., & Thorpe, S. (2001). A limit to the speed of processing in ultra-rapid visual categorization of novel natural scenes. *Journal of Cognitive Neuroscience, 13,* 171-180.
Felleman, D. J. & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex, 1,* 1-47.

Franks, J. J. & Bransford, J. D. (1971). Abstraction of Visual Patterns. *Journal of Experimental Psychology, 90,* 65-&.
Frith, U. (1974). Curious Effect with Reversed Letters Explained by A Theory of Schema. *Perception & Psychophysics, 16,* 113-116.
Fulkerson, A. L. & Haaf, R. A. (2003). The influence of labels, non-labeling sounds, and source of auditory input on 9-and 15-month-olds' object categorization. *Infancy, 4,* 349-369.
Fulkerson, A. L. & Haaf, R. A. (2006). Does object naming aid 12-month-olds' formation of novel categories? *First Language, 26,* 347-361.
Garner, W. R. & Clement, D. E. (1963). Goodness of Pattern and Pattern Uncertainty. *Journal of Verbal Learning and Verbal Behavior, 2,* 446-452.
Gauthier, I., James, T. W., Curby, K. M., & Tarr, M. J. (2003). The influence of conceptual knowledge on visual discrimination. *Cognitive Neuropsychology, 20,* 507-523.
Gelman, S. A. & Coley, J. D. (1991). Language and categorization: The acquisition of natural kind terms. In J.P.Byrnes & S. A. Gelman (Eds.), *Perspectives on language and thought: Interrelations in development* (pp. 146-196). Cambridge: Cambrudge University Press.
Gentner, D. & Goldin-Meadow, S. (2003). *Language in Mind: Advances in the Study of Language and Thought*. Cambridge, MA.: MIT Press.
Gentner, D. & Loewenstein, J. (2002). Relational language and relational thought. In J.Byrnes & E. Amsel (Eds.), *Language, Literacy, and Cognitive Development* (pp. 87-120). Mahwah, NJ: LEA.
Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S.Kuczaj (Ed.), *Language development: Language, cognition, and culture* (Hillsdale, NJ: Erlbaum.
Gibson, B. S. & Kingstone, A. (2006). Visual attention and the semantics of space: Beyond central and peripheral cues. *Psychological Science, 17,* 622-627.
Gibson, E. J. (1953). Improvement in Perceptual Judgments As A Function of Controlled Practice Or Training. *Psychological Bulletin, 50,* 401-431.
Gibson, J. J. & Gibson, E. J. (1955). Perceptual Learning - Differentiation Or Enrichment. *Psychological Review, 62,* 32-41.
Gilbert, A. L., Regier, T., Kay, P., & Ivry, R. B. (2006). Whorf hypothesis is supported in the right visual field but not the left. *Proceedings of the National Academy of Sciences of the United States of America, 103,* 489-494.
Gleitman, L. & Papafragou, A. (2005). Language and thought. In K.Holyoak & B. Morrison (Eds.), *Cambridge Handbook of thinking and Reasoning* (pp. 633-661). Cambridge: Cambridge University Press.
Goldstein, K. (1924). Das Wesen der amnestischen Aphasie. *Schweizer Archiv fuËr Neurologieund Psychiatrie, 19,* 163-175.
Goldstein, K. (1948). *Language and language disturbances*. New York: Grune & Stratton.
Goldstone, R. (1994). Influences of Categorization on Perceptual Discrimination. *Journal of Experimental Psychology-General, 123,* 178-200.
Goldstone, R. L., Steyvers, M., & Larimer, K. (1996). Categorical perception of novel dimensions. In (pp. 243-248). Hillsdale, NJ: Lawrence Erlbaum Associates.
Goldstone, R. L., Steyvers, M., Spencer-Smith, J., & Kersten, A. W. (2004). Interactions Between Perceptual and Conceptual Learning. In E.Diettrich & A. B. Markman (Eds.), *Cognitive Dynamics: Conceptual Change in Humans and Machines* (pp. 191-228). Mahwah, NJ: Lawrence Erlbaum and Associates.
Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology, 49,* 585-612.
Goldstone, R. L. & Barsalou, L. W. (1998). Reuniting perception and conception. *Cognition, 65,* 231-262.
Goldstone, R. L., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations through category learning. *Cognition, 78,* 27-43.
Goldstone, R. L., Steyvers, M., & Rogosky, B. J. (2003). Conceptual interrelatedness and caricatures. *Memory & Cognition, 31,* 169-180.
Goss, A. E. (1953). Transfer As A Function of Type and Amount of Preliminary Experience with Task Stimuli. *Journal of Experimental Psychology, 46,* 419-427.
Goss, A. E. (1961). Verbal Mediating Responses and Concept-Formation. *Psychological Review, 68,* 248-274.
Grill-Spector, K. & Kanwisher, N. (2005). Visual recognition - As soon as you know it is there, you know what it is. *Psychological Science, 16,* 152-160.
Grossman, M. (1978). Game of Name - Examination of Linguistic Reference After Brain-Damage. *Brain and Language, 6,* 112-119.
Gumperz, J. J. & Levinson, S. C. (1996). *Rethinking linguistic relativity*. Cambridge: Cambridge University Press.

Hake, H. W. & Eriksen, C. W. (1956). Role of Response Variables in Recognition and Identification of Complex Visual Forms. *Journal of Experimental Psychology, 52,* 235-243.
Harnad, S. (1987). *Categorical Perception: The Groundwork of Cognition*. New York: Cambridge University Press.
Harnad, S. (2005). Cognition is categorization. In H.Cohen & C. Lefebvre (Eds.), *Handbook of Categorization* (Elsevier.
Heaton, P., Ludlow, A., & Roberson, D. When less is more: poor discrimination but good colour memory in Autism. *Autism Spectrum Disorders,* (in press).
Hermer-Vazquez, L., Spelke, E. S., & Katsnelson, A. S. (1999). Sources of flexibility in human cognition: Dual-task studies of space and language. *Cognitive Psychology, 39,* 3-36.
Herrnstein, R. J. (1990). Levels of Stimulus-Control - A Functional-Approach. *Cognition, 37,* 133-166.
Hespos, S. J. & Spelke, E. S. (2004). Conceptual precursors to language. *Nature, 430,* 453-456.
Hochstein, S. & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron, 36,* 791-804.
Hommel, B., Pratt, J., Colzato, L., & Godijn, R. (2001). Symbolic control of visual attention. *Psychological Science, 12,* 360-365.
Hupe, J. M., James, A. C., Girard, P., Lomber, S. G., Payne, B. R., & Bullier, J. (2001). Feedback connections act on the early part of the responses in monkey visual cortex. *Journal of Neurophysiology, 85,* 134-145.
James, W. (1890). *Principles of Psychology. Vol. 1*. New York: Holt.
Johnston, J. R. (1981). On location: Thinking and talking about space. *Topics in Language Disorders, 2,* 17-32.
Johnston, J. R. (1985). Cognitive prerequisites: The evidence from children learning English. In D.I.Slobin (Ed.), *The cross-linguistic study of language acquisition, Vol. 1: Theoretical issues* (pp. 961-1004). Hillsdale, NJ: Erlbaum.
Johnston, J. R. (1988). Children's verbal representation of spatial location. In J.Stiles, M. Kritchevsky, & U. Bellugi (Eds.), *Spatial Cognition: Brain Bases and Development* (Hillsdale, NJ: Erlbaum.
Johnston, W. A., Hawley, K. J., & Farnham, J. M. (1993). Novel Popout - Empirical Boundaries and Tentative Theory. *Journal of Experimental Psychology-Human Perception and Performance, 19,* 140-153.
Jonides, J. & Gleitman, H. (1972). Conceptual Category Effect in Visual Search - O As Letter Or As Digit. *Perception & Psychophysics, 12,* 457-460.
Kail, R. & Nippold, M. A. (1984). Unconstrained Retrieval from Semantic Memory. *Child Development, 55,* 944-951.
Karlin, M. B. & Bower, G. H. (1976). Semantic Category Effects in Visual Word Search. *Perception & Psychophysics, 19,* 417-424.
Katz, P. A. (1963). Effects of Labels on Children's Perception and Discrimination-Learning. *Journal of Experimental Psychology, 66,* 423-&.
Katz, P. A., Karp, B., & Yalisove, D. (1970). Verbal Mediation of Childrens Perception - Role of Response Variables. *Journal of Experimental Psychology, 85,* 349-&.
Katz, P. A. & Zigler, E. (1969). Effects of Labels on Perceptual Transfer - Stimulus and Developmental Factors. *Journal of Experimental Psychology, 80,* 73-&.
Kelter, S., Cohen, R., Engel, D., List, G., & Strohner, H. (1976). Aphasic Disorders in Matching Tasks Involving Conceptual Analysis and Covert Naming. *Cortex, 12,* 383-394.
Kelter, S., Cohen, R., Engel, D., List, G., & Strohner, H. (1977). Conceptual Structure of Aphasic and Schizophrenic-Patients in A Nonverbal Sorting Task. *Journal of Psycholinguistic Research, 6,* 279-303.
Koemeda-Lutz, M., Cohen, R., & Meier, E. (1987). Organization of and Access to Semantic Memory in Aphasia. *Brain and Language, 30,* 321-337.
Kosslyn, S. M., Ganis, G., & Thompson, W. L. (2001). Neural foundations of imagery. *Nature Reviews Neuroscience, 2,* 635-642.
Kosslyn, S. M., Pascual-Leone, A., Felician, O., Camposano, S., Keenan, J. P., Thompson, W. L. et al. (1999). The role of Area 17 in visual imagery: Convergent evidence from PET and rTMS. *Science, 284,* 167-170.
Kotovsky, L. & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development, 67,* 2797-2822.
Krueger, L. E. (1984). The Category Effect in Visual-Search Depends on Physical Rather Than Conceptual Differences. *Perception & Psychophysics, 35,* 558-564.
Kuhl, P. (1994). Learning and representation in speech and language. *Current Opinions in Neurobiology, 4,* 812-822.

Lamme, V. A. F., Rodriguez-Rodriguez, V., & Spekreijse, H. (1999). Separate processing dynamics for texture elements, boundaries and surfaces in primary visual cortex of the macaque monkey. *Cerebral Cortex, 9,* 406-413.
Lamme, V. A. F. & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences, 23,* 571-579.
Lamme, V. A. F., Super, H., Landman, R., Roelfsema, P. R., & Spekreijse, H. (2000). The role of primary visual cortex (V1) in visual awareness. *Vision Research, 40,* 1507-1521.
Lamme, V. A. F., Super, H., & Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology, 8,* 529-535.
Landau, B. & Shipley, E. (2001). Labelling patterns and object naming. *Developmental Science, 4,* 109-118.
Lawrence, D. H. (1949). Acquired Distinctiveness of Cues .1. Transfer Between Discriminations on the Basis of Familiarity with the Stimulus. *Journal of Experimental Psychology, 39,* 770-784.
Lawrence, D. H. (1950). Acquired Distinctiveness of Cues .2. Selective Association in A Constant Stimulus Situation. *Journal of Experimental Psychology, 40,* 175-188.
Lee, T. S. & Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences of the United States of America, 98,* 1907-1911.
Lemer, C., Dehaene, S., Spelke, E., & Cohen, L. (2003). Approximate quantities and exact number words: dissociable systems. *Neuropsychologia, 41,* 1942-1958.
Levin, D. T. & Angelone, B. L. (2001). Visual search for a socially defined feature: What causes the search asymmetry favoring cross-race faces? *Perception & Psychophysics, 63,* 423-435.
Levin, D. T., Takarae, Y., Miner, A. G., & Keil, F. (2001). Efficient visual search by category: Specifying the features that mark the difference between artifacts and animals in preattentive vision. *Perception & Psychophysics, 63,* 676-697.
Levinson, S. C. (2003). *Space in Language and Cognition*. Cambridge: Cambridge University Press.
Levinson, S. C., Kita, S., Haun, D. B. M., & Rasch, B. H. (2002). Returning the tables: language affects spatial reasoning. *Cognition, 84,* 155-188.
Li, P. & Gleitman, L. (2002). Turning the tables: language and spatial reasoning. *Cognition, 83,* 265-294.
Liberman, A. M., Cooper, F. S., Shankwei, D. P., & Tuddert, M. (1967). Perception of Speech Code. *Psychological Review, 74,* 431-&.
Loftus, G. R. & Masson, M. E. J. (1994). Using Confidence-Intervals in Within-Subject Designs. *Psychonomic Bulletin & Review, 1,* 476-490.
Logan, G. D. (1995). Linguistic and Conceptual Control of Visual-Spatial Attention. *Cognitive Psychology, 28,* 103-174.
Lucy, J. A. (1992). *Language Diversity and Thought : A Reformulation of the Linguistic Relativity Hypothesis*. Cambridge, MA: Cambridge University Press.
Lupyan, G. (2007, in press). From Chair To "Chair:" A Representational Shift Account Of Object Labeling Effects On Memory. *Journal of Experimental Psychology: General* .
Lupyan, G. (2005). When Naming Means Forgetting: Verbal Classification Leads to Worse Memory. In *Twenty-Seventh Annual Meeting of the Cognitive Science Society* (pp. 2518). Mahwah, NJ: Lawrence Erlbaum.
Lupyan, G. (2007b). Reuniting Categories, Language, and Perception. In *Twenty-Ninth Annual Meeting of the Cognitive Science Society*.
Lupyan, G., Rakison, D. H., & McClelland, J. L. Language is not just for talking: labels facilitate learning of novel categories. *Psychological Science,* (in press).
Macknik, S. L. & Martinez-Conde, S. (2007). The role of feedback in visual masking and visual processing. *Advances in Cognitive Psychology, 3,* 125-152.
Malinowski, P. & Hubner, R. (2001). The effect of familiarity on visual-search performance: Evidence for learned basic features. *Perception & Psychophysics, 63,* 458-463.
Mandler, J. M. (1996). Preverbal representation and language. In P.Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and Space* (Cambridge, MA: MIT Press.
Maunsell, J. H. R. & Newsome, W. T. (1987). Visual Processing in Monkey Extrastriate Cortex. *Annual Review of Neuroscience, 10,* 363-401.
Mccleary, C. (1988). The Semantic Organization and Classification of 14 Words by Aphasic Patients. *Brain and Language, 34,* 183-202.
McClelland, J. L. & Chappell, M. (1998). Familiarity breeds differentiation: A subjective-likelihood approach to the effects of experience in recognition memory. *Psychological Review, 105,* 724-760.
McClelland, J. L. & Elman, J. L. (1986). The Trace Model of Speech-Perception. *Cognitive Psychology, 18,* 1-86.

Mccune-Nicolich, L. (1981). The Cognitive Bases of Relational Words in the Single Word Period. *Journal of Child Language, 8,* 15-34.
McDonough, L., Choi, S., & Mandler, J. M. (2003). Understanding spatial relations: Flexible infants, lexical adults. *Cognitive Psychology, 46,* 229-259.
Mckinney, J. P. (1963). Disappearance of Luminous Designs. *Science, 140,* 403-&.
Mesulam, M. M. (1998). From sensation to cognition. *Brain, 121,* 1013-1052.
Miller, N. E. & Dollard, J. (1941). *Social Learning and Imitation*. New Haven: Yale University Press.
Moran, J. & Desimone, R. (1985). Selective Attention Gates Visual Processing in the Extrastriate Cortex. *Science, 229,* 782-784.
Nazzi, T. & Gopnik, A. (2001). Linguistic and cognitive abilities in infancy: when does language become a tool for categorization? *Cognition, 80,* B11-B20.
Nelson, K. (1974). Concept, Word, and Sentence - Interrelations in Acquisition and Development. *Psychological Review, 81,* 267-285.
Noppeney, U. & Wallesch, C. W. (2000). Language and cognition - Kurt Goldstein's theory of semantics. *Brain and Cognition, 44,* 367-386.
Oakes, L. M. & Madole, K. L. (2000b). The future of infant categorization research: A process-oriented approach. *Child Development, 71,* 119-126.
Oakes, L. M. & Madole, K. L. (2000a). The future of infant categorization research: A process-oriented approach. *Child Development, 71,* 119-126.
Ozgen, E. (2004). Language, Learning, and Color Perception. *Current Directions in Psychological Science, 13,* 95-98.
Ozgen, E. & Davies, I. R. L. (2002). Acquisition of categorical color perception: A perceptual learning approach to the linguistic relativity hypothesis. *Journal of Experimental Psychology-General, 131,* 477-493.
Pettet, M. W. & Gilbert, C. D. (1992). Dynamic Changes in Receptive-Field Size in Cat Primary Visual-Cortex. *Proceedings of the National Academy of Sciences of the United States of America, 89,* 8366-8370.
Pica, P., Lemer, C., Izard, W., & Dehaene, S. (2004). Exact and approximate arithmetic in an Amazonian indigene group. *Science, 306,* 499-503.
Pilling, M., Wiggett, A., Ozgen, E., & Davies, I. R. L. (2003). Is color "categorical perception" really perceptual? *Memory & Cognition, 31,* 538-551.
Pinker, S. (1995). *The Language Instinct*. New York: Harper Collins.
Plaut, D. C. (2002). Graded modality-specific specialisation in semantics: A computational account of optic aphasia. *Cognitive Neuropsychology, 19,* 603-639.
Plunkett, K., Hu, J.-F., & Cohen, L. B. (2007). Labels can override perceptual categories. *Cognition, doi:10.1016/j.cognition.2007.04.003*.
Posner, M. I. (1970). On Relationship Between Letter Names and Superordinate Categories. *Quarterly Journal of Experimental Psychology, 22,* 279-&.
Posner, M. I. & Keele, S. W. (1968). On Genesis of Abstract Ideas. *Journal of Experimental Psychology, 77,* 353-&.
Posner, M. I. & Mitchell, R. F. (1967). Chronometric Analysis of Classification. *Psychological Review, 74,* 392-&.
Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences, 22,* 341-+.
Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature, 435,* 1102-1107.
Ratterman, M. J. & Gentner, D. (1998). The effect of language on similarity: The use of relational symbols improves young children's performance on a mapping task. In K.Holyoak, D. Gentner, & B. Kokinov (Eds.), *Advances in analogy research: Integration of theory and data from the cognitive, computational and neural sciences.* (pp. 274-281). Sophia: New Bulgarian University.
Rauschenberger, R. & Yantis, S. (2006). Perceptual encoding efficiency in visual search. *Journal of Experimental Psychology-General, 135,* 116-131.
Reali, F., Spivey, M., Tyler, M., & Terranova, J. (2006). Inefficient conjunction search made efficient by concurrent spoken delivery of target identity. *Perception and Psychophysics, in press*.
Rensink, R. A. (2000). The dynamic representation of scenes. *Visual Cognition, 7,* 17-42.
Richards, J. T. (1978). Interitem Structure and Facilitation of Simultaneous Comparison. *Journal of Experimental Psychology-Human Perception and Performance, 4,* 72-87.
Richards, J. T. & Reicher, G. M. (1978). Effect of Background Familiarity in Visual-Search - Analysis of Underlying Factors. *Perception & Psychophysics, 23,* 499-505.

Roberson, D. & Davidoff, J. (2000). The categorical perception of colors and facial expressions: The effect of verbal interference. *Memory & Cognition, 28,* 977-986.
Roberson, D., Davidoff, J., & Braisby, N. (1999). Similarity and categorisation: neuropsychological evidence for a dissociation in explicit categorisation tasks. *Cognition, 71,* 1-42.
Roberson, D., Davidoff, J., Davies, I. R. L., & Shapiro, L. R. (2005). Color categories: Evidence for the cultural relativity hypothesis. *Cognitive Psychology, 50,* 378-411.
Robinson, J. S. (1955). The Effect of Learning Verbal Labels for Stimuli on Their Later Discrimination. *Journal of Experimental Psychology, 49,* 112-114.
Roelfsema, P. R., Lamme, V. A. F., & Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature, 395,* 376-381.
Rogers, T. T. & McClelland, J. L. (2004). *Semantic Cognition: A Parallel Distributed Processing Approach*. Cambridge, MA: Bradford Book.
Rosch, E. (1973). On the Internal Structure of Perceptual and Semantic Categories. In T.E.Moore (Ed.), *Cognitive Development and the Acquisition of Language* (New York: Academic Press.
Rosch, E. (1978). Principles of Categorization. In E.Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27-48). Hillsdale, NJ: Erlbaum.
Rossion, B. & Pourtois, G. (2004). Revisiting Snodgrass and Vanderwart's object pictorial set: The role of surface detail in basic-level object recognition. *Perception, 33,* 217-236.
Rossman, I. L. & Goss, A. E. (1951). The Acquired Distinctiveness of Cues - the Role of Discriminative Verbal Responses in Facilitating the Acquisition of Discriminative Motor Responses. *Journal of Experimental Psychology, 42,* 173-182.
Rumelhart, D. E. & McClelland, J. L. (1982). An Interactive Activation Model of Context Effects in Letter Perception .2. the Contextual Enhancement Effect and Some Tests and Extensions of the Model. *Psychological Review, 89,* 60-94.
Rumelhart, D. E., McClelland, J. L., & the PDP Research Group (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volumes 1 and 2*. Cambridge, MA: MIT Press.
Rumelhart, D. E., Smolensky, D. E., McClelland, J. L., & Hinton, G. E. (1986). Parallel Distributed Processing Models of Schemata and Sequential Thought Processes. In J.L.McClelland & D. E. Rumelhart (Eds.), *Parallel Distributed Processing Vol II* (pp. 7-57). Cambridge, MA: MIT Press.
Schlesinger, I. M. (1977). Role of Cognitive-Development and Linguistic Input in Language-Acquisition. *Journal of Child Language, 4,* 153-169.
Serences, J. T. & Yantis, S. (2006). Selective visual attention and perceptual coherence. *Trends in Cognitive Sciences, 10,* 38-45.
Shen, J. & Reingold, E. M. (2007). Visual search asymmetry: The influence of stimulus familiarity and low-level features. *Perception & Psychophysics, 63,* 464-475.
Shiffrin, R. M. & Steyvers, M. (1997). Model for recognition memory: REM - Retrieving effectively from memory. *Psychonomic Bulletin & Review, 4,* 145-166.
Slobin, D. I. (1996). From "Thought and Language" to "Thinking for Speaking". In J.J.Gumperz & S. C. Levinson (Eds.), *Rethinking linguistic relativity* (pp. 70-96). Cambridge: Cambridge University Press.
Slobin, D. I. (1985). Crosslinguistic evidence for the language-making capacity. In D.I.Slobin (Ed.), *The cross-linguistic study of language acquisition, Vol. 1: Theoretical issues* (Hillsdale, NJ: Erlbaum.
Sloutsky, V. M. & Fisher, A. V. (2004a). Induction and categorization in young children: A similarity-based model. *Journal of Experimental Psychology-General, 133,* 166-188.
Sloutsky, V. M. & Fisher, A. V. (2004b). When development and learning decrease memory - Evidence against category-based induction in children. *Psychological Science, 15,* 553-558.
Smilek, D., Dixon, M. J., & Merikle, P. M. (2006). Revisiting the category effect: The influence of meaning and search strategy on the efficiency of visual search. *Brain Research, 1080,* 73-90.
Smith, L. B., Colunga, E., & Yoshida, H. (2003). Making an Ontology: Cross-linguistic evidence. In D.H.Rakison & L. M. Oakes (Eds.), *Early Category and Concept Development: Making Sense of the Blooming, Buzzing Confusion* (pp. 275-302). Oxford University Press.
Spelke, E. S. & Tsivkin, S. (2001). Language and number: a bilingual training study. *Cognition, 78,* 45-88.
Spiker, C. C. (1956). Stimulus Pretraining and Subsequent Performance in the Delayed-Reaction Experiment. *Journal of Experimental Psychology, 52,* 107-111.
Spivey, M. J. (2007). *The Continuity of Mind*. Oxford University Press.
Spivey, M. J. & Dale, R. (2004). On the continuity of mind: Toward a dynamical account of cognition. In B.Ross (Ed.), *The Psychology of Learning and Motivation Vol. 45* (Elsevier.

Spivey, M. J. & Dale, R. (2005). On the continuity of mind: Toward a dynamical account of cognition. In B. Ross (Ed.), *The Psychology of Learning and Motivation Vol. 45* (pp. 85-142). Elsevier.
Spivey, M. J., Tyler, M. J., Eberhard, K. M., & Tanenhaus, M. K. (2001). Linguistically mediated visual search. *Psychological Science, 12,* 282-286.
Spivey-Knowlton, M. J., Trueswell, J. C., & Tanenhaus, M. K. (1993). Context Effects in Syntactic Ambiguity Resolution - Discourse and Semantic Influences in Parsing Reduced Relative Clauses. *Canadian Journal of Experimental Psychology-Revue Canadienne de Psychologie Experimentale, 47,* 276-309.
Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of Visual and Linguistic Information in Spoken Language Comprehension. *Science, 268,* 1632-1634.
Taylor, T. L. & Hamm, J. (1997). Category effects in temporal visual search. *Canadian Journal of Experimental Psychology-Revue Canadienne de Psychologie Experimentale, 51,* 36-46.
Theeuwes, J. (1993). Visual Selective Attention - A Theoretical-Analysis. *Acta Psychologica, 83,* 93-154.
Tighe, L. S. & Tighe, T. J. (1966). Discrimination Learning - 2 Views in Historical Perspective. *Psychological Bulletin, 66,* 353-&.
Tovee, M. J. (1994). How Fast is the Speed of Thought? *Current Biology, 4,* 1125-1127.
Townsend, J. T. (1990). Serial Vs Parallel Processing - Sometimes They Look Like Tweedledum and Tweedledee But They Can (and Should) be Distinguished. *Psychological Science, 1,* 46-54.
Treisman, A. & Gormican, S. (1988). Feature Analysis in Early Vision - Evidence from Search Asymmetries. *Psychological Review, 95,* 15-48.
Treisman, A. M. & Gelade, G. (1980). Feature-Integration Theory of Attention. *Cognitive Psychology, 12,* 97-136.
Treue, S. (2003). Climbing the cortical ladder from sensation to perception. *Trends in Cognitive Sciences, 7,* 469-471.
Vernon, M. D. (1955). The Functions of Schemata in Perceiving. *Psychological Review, 62,* 180-192.
Vidyasagar, T. R. (1999). A neuronal model of attentional spotlight: parietal guiding the temporal. *Brain Research Reviews, 30,* 66-76.
Vignolo, L. A. (1999). Disorders of conceptual thinking in aphasia. In G.Denes & L. Pizzamiglio (Eds.), *Handbook of Clinical and Experimental Neuropsychology* (pp. 273-288). Hove: Psychology Press.
Vygotsky, L. S. (1962). *Thought and Language*. Cambridge, MA: MIT Press.
Wang, Q. Q., Cavanagh, P., & Green, M. (1994). Familiarity and Pop-Out in Visual-Search. *Perception & Psychophysics, 56,* 495-500.
Waxman, S. (2004). Everything had a name, and each name gave birth to a new thought: Links between early word-learning and conceptual organization. In D.G.Hall & S. Waxman (Eds.), (Cambridge, MA: MIT Press.
Waxman, S. (1999). The dubbing ceremony revisited: Object naming and categorization in infancy and early childhood. In D.L.Medin & S. Atran (Eds.), *Folkbiology* (pp. 233-284). Cambridge, MA: MIT Press.
Waxman, S. & Booth, A. (2003). The origins and evolution of links between word learning and conceptual organization: new evidence from 11-month-olds. *Developmental Science, 6,* 128-135.
Waxman, S. R. & Hall, D. G. (1993). The Development of A Linkage Between Count Nouns and Object Categories - Evidence from 15-Month-Old to 21-Month-Old Infants. *Child Development, 64,* 1224-1241.
Waxman, S. R. & Markow, D. B. (1995). Words as invitations to form categories: Evidence from 12- to 13-month-old infants. *Cognitive Psychology, 29,* 257-302.
Welford, T. A. (1980). *Reaction Times*. New York: Academic Press.
White, M. J. (1977). Identification and Categorization in Visual-Search. *Memory & Cognition, 5,* 648-657.
Whorf, B. L. (1956). *Language, Thought, and Reality*. Cambridge, MA: MIT Press.
Williams, Z. M., Elfar, J. C., Eskandar, E. N., Toth, L. J., & Assad, J. A. (2003). Parietal activity and the perceived direction of ambiguous apparent motion. *Nature Neuroscience, 6,* 616-623.
Winawer, J., Witthoft, N., Frank, M. C., Wu, L., Wade, A. R., & Boroditsky, L. (2007). Russian blues reveal effects of language on color discrimination. *Proceedings of the National Academy of Sciences of the United States of America, 104,* 7780-7785.
Wittgenstein, L. (1953). *Philosophical investigations*. Oxford: Basil Blackwell.
Wolfe, J. M. (1994). Guided Search 2.0 - A Revised Model of Visual-Search. *Psychonomic Bulletin & Review, 1,* 202-238.
Wolfe, J. M. (1998). What can 1 million trials tell us about visual search? *Psychological Science, 9,* 33-39.
Wolfe, J. M. (2001). Asymmetries in visual search: An introduction. *Perception & Psychophysics, 63,* 381-389.
Wolfe, J. M. (2003). Moving towards solutions to some enduring controversies in visual search. *Trends in Cognitive Sciences, 7,* 70-76.

Wolfe, J. M. & Bennett, S. C. (1997). Preattentive object files: Shapeless bundles of basic features. *Vision Research, 37,* 25-43.
Wolfe, J. M., Butcher, S. J., Lee, C., & Hyle, M. (2003). Changing your mind: On the contributions of top-down and bottom-up guidance in visual search for feature singletons. *Journal of Experimental Psychology-Human Perception and Performance, 29,* 483-502.
Wolfe, J. M. & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience, 5,* 495-501.
Wolfe, J. M., Horowitz, T. S., Kenner, N., Hyle, M., & Vasan, N. (2004). How fast can you change your mind? The speed of top-down guidance in visual search. *Vision Research, 44,* 1411-1426.
Wolfe, J. M., Stewart, M. I., Friedmanhill, S. R., & Oconnell, K. M. (1992). The Role of Categorization in Visual-Search for Orientation. *Journal of Experimental Psychology-Human Perception and Performance, 18,* 34-49.
Xu, F. (1999). Object individuation and object identity in infancy: The role of spatiotemporal information, object property information, and language. *Acta Psychologica, 102,* 113-136.
Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. *Cognition, 85,* 223-250.
Yoshida, H. & Smith, L. B. (2005). Linguistic cues enhance the learning of perceptual cues. *Psychological Science, 16,* 90-95.
Younger, B. (1990). Infant Categorization - Memory for Category-Level and Specific Item Information. *Journal of Experimental Child Psychology, 50,* 131-155.
Younger, B. A. & Cohen, L. B. (1985). How Infants Form Categories. *Psychology of Learning and Motivation-Advances in Research and Theory, 19,* 211-247.
Zurif, E. B., Caramazza, Myerson, R., & Galvin, J. (1974). Semantic Feature Representations for Normal and Aphasic Language. *Brain and Language, 1,* 167-187.