On Features

J. Kevin Varden

Abstract

To learn a language, a child must make sense of a huge amount of incoming stimuli and correctly apply labels to objects in their field of sensory input. To facilitate the parsing of the speech stream, the child is thought to use specific prosodic and articulatory features; to facilitate the organization of objects discerned in that aural stimuli, it is thought that the child uses semantic features—traits and attributes of objects in our sensory environment that are salient to even the most naïve language learner. Semantic features—that is, discernible attributes of objects—allow us to group objects according to traits and trait sets, and distinguish them according to disjoint trait sets. To facilitate investigations into word meaning based on or containing a component feature-checking algorithm, a database of semantic features of concrete objects, spatiality and time, called cogPrime, has been constructed (see Varden 2001 for details).

This paper discusses the relationship between 'reality', word meaning, and feature use, and reiterates the position that an apparent overlap in articulatory, auditory and semantic feature organization and use can tell us much about how we categorize sensory input. It also discusses an alternate way of eliciting semantic features other than the use of printed natural language.

1 Introduction

One thing has become clear in the present investigation: the use of features (i.e. salient attributes of objects) has a place in all aspects of cognition. One cannot discuss feature use without making reference to *categorization*, the parsing and sorting of incoming stimuli into various categories or classes (e.g. voiced vs. voiced segment; object vs. background, animal vs. tool, etc.). It therefore seems worthwhile to bring together some of the research conducted in the various disciplines with reference to the nature of features, their organization and use, and their relation to categorization, with an eye on future integration of the parts into a whole.

The organization of the paper is as follows. §2 makes several points about the nature of the reality that we are categorizing by extracting features from and assigning features to. §3 discusses the nature of word meaning as seen from various perspectives, including the use of adjectives as feature assignors in the current line of research. §4 takes a look at various aspects of features, including their expression and their use, while §5 presents an alternative to the use of natural language to elicit features associated with object attributes. Finally, §6 summarizes the main points of the paper.

2 On perceiving 'reality'

As a beginning, it seems useful to briefly review exactly what the context for a discussion of features is. This section briefly reviews the most basic of all perceptual activities involving features: categorization.

2.1 Learning is processing and making judgments about 'reality'

In order to learn to speak, a child must make sense of an enormous amount of input. Not only are all of the operating senses inundated with wave after wave of sensory stimuli, at a bare minimum the child must be able to discriminate the incoming aural language input from the background, and determine which object(s) is being referred to by the speaker to learn the correct labels for objects. Among the processes the child must call into action are discrimination of sounds, categorization, segmentation, attention to referenced items, mental representation and memory (Jusczyk 1992, 1994).

Much of learning about the world seems to involve constructing a 'world' in our head (e.g. Jackendoff 1983; see Varden 2001 for an overview). This is supported by, in addition to other things, the fact that incoming stimuli is in constant flux, particularly so for visual data (e.g. Seung & Lee 2000). And yet we perceive non-moving objects as stationary. Thus we constantly ignore sensory fluctuations and see what we 'know' is true: stationary objects are constant and stationary, not constantly fluctuating.

The normal fluctuation of the incoming stimuli reflecting objects in our environment does not adversely affect the most basic of processes utilizing that stimuli: categorization, the assigning of incoming stimuli to a certain category or class for the purposes of identification.

2.2 Categorization in language acquisition

Children apparently categorize the objects they are learning labels for as a means to facilitate their learning. There appear to be 3 generalizations used by children in learning the labels for 3 objects (i.e. in categorizing those new objects): the principle of *whole object assumption,* whereby a new word is assumed to refer to a whole object as opposed to a part or substance; the *taxonomic assumption*, whereby an object's label is extended to other object's of the same category; and the *mutual exclusivity assumption*, whereby each object will be assigned one and only one label (see Markman 1994 and references therein). The whole object assumption would seem to guide children's early language acquisition; the others help them to acquire other necessary labels. These distinctions are critical for helping the child break into their language's semantic code.[1]

This happens very early in development. Landau (2000: 214*ff)* reports that children as young as two, as well as adults, crucially rely on shape in judging objects; older children and adults will do so if simply naming objects, but can be more flexible in their categorizations depending on the task at hand. Needham (1999; 2001) presents evidence that infants as young as 4.5 mos. can individuate objects in their environment based on features—they attend to shape much more than color and pattern to judge objects' boundaries, and use attributes of objects to determine whether or not an object has changed for moved, much in the same way that adults do.

Indeed, categorization is such a prevalent and basic activity that it has even been suggested to be the source of words—that iconic and categorical representations interacted to yield a third, the symbolic representations that are at the heart of human language (Harnad 1996).

2.3 Categorization is highly structured

However, we do not simply maintain a hodge-podge collection of mental objects and relationships. Rosch (et al. 1976) provides solid evidence for the use of basic categories in interpreting our environment. Based on English nouns, they identified four attributes of members of the most inclusive categories: they have a significant number of attributes in common, move and change in roughly the same way, have similar shapes, and can be, at least roughly, identified by exemplars from their categories.

In addition to these basic categories, further divisions can be made due to the linguistic experience and needs of the language learner. A good example is the abstract food categories like *snacks* that people make use of (Ross & Murphy 1999); such categories are dependent on individual experience, and yet become a dominant force in

---

[1] But see Mohanan 1992 for an alternative view of language as a self-developing system, not a 'code-breaking' activity.

the mature speaker (Rehder & Ross 2001). Development of numeral classifiers systems in various languages (e.g. Matsumoto 1993 for Japanese) point to experience-derived classifications as well.

This experience evidently leads to a highly structured, rich system for classifying objects according to incoming stimuli. (The flexibility of that system in the mature speaker will be discussed in §4.7.)

2.4 Categorization is a higher-order function

Another point that bears stressing is that the use our brain makes of the auditory features that strike our ears may be quite like the use it makes of the visual stimuli that strikes our eyes, or the tactile information being presented to our fingers (and in the case of infants, mouths). Each modality is, of course, quite different, and both the visual and aural processing pathways are distinct. However, it has also been indicated that the higher-order processing of stimuli (i.e. cognition) is independent of the lower-order processing (i.e. perception)[2]. In regard to aural stimuli, the cochlea in the inner ear performs a basic Fourier analysis device for the initial identification of the pitch of the sound; the neural impulses produced by this transformation are then presented (along with other information) to the auditory nerve (Borden et al. 1994). This initial processing is quite independent of whether or not later processing centers identify the sound as speech or not (e.g. Gokcen & Fox 2001).

This independence of the higher-order processing extends to whether or not we pay attention to what is plainly speech. Someone we love—or more so, someone we want to love—triggers high attention; multiple voices in a crowded room or our parents scolding us trigger low or non-existent attention. Likewise, the visual processing that occurs at the low level (i.e. the neural response to the light striking the back of the eye) is quite separate from the later identification of objects, movement, and categorization of objects. No matter what the modality of the stimuli, it can be perceived by the senses without the person being aware of the perception (for an overview of the relevant research, see Merikle et al. 2001).

Thomas et al. (2001) provided good evidence that the neural cells involved in

---

[2] Although perception and cognition do influence each other; see, e.g., Goldstone & Barsalou (1998) for an overview.

categorizing objects (in their case, macaque monkeys categorizing fish and trees) was handled by cells that are not category specific; the categorization occurred over the sum of the nerves involved as a higher-order function. Hinojosa et al. (2001) provide good evidence that, although some differences in the neural activity can be seen, the same areas of the brain handle the categorization of very different objects (animals and tools).

If this can be extended to other modalities, and higher-order functions are responsible for both visual categorization and speech categorization, they both might well involve the same general higher-order principles: in simplistic terms, something on the order of ‘take these stimuli and sort them into groups’. If this is so, than the use our brain makes of discerning attributes of our environment might to well to borrow the mechanisms that afford phonology and phonetics much of their explanatory power, specifically, underspecification of features, feature bundling and hierarchical organization of attributes being attended to in the manner of Feature Geometry. This has been done to some degree in the psychological research, even if the same labels have not been applied (e.g. the category hierarchies in Schyns 1998; Lundqvist et al. 1999; Ahn et al. 2000). Future research will surely benefit from a more extensive comparison of the processing of the two—and all other—sensory modalities. This is in line with verifying and validating an overall perceptual system such as that called for in Kuhl (1987); perhaps the Adaptation-Level theory outlined in Wilson (1987), or the connectionist models of Kruschke (1990, 1996); Kruschke & Johansen (1999); and Kalish & Kruschke (2000).

This is not to say that incoming stimuli are simply assigned a label which persists while higher-level processing occurs. Perception and cognition are more closely inter-related than that.

2.5 Categorization is closely tied to perception

The objects that we perceive are categorized not just by incoming perceptions, but also according to our expectations. Oliva & Schyns (1997, 2000) and Schyns & Oliva (1999) have shown that the range of features attended to when categorizing objects or people is influenced by the features the subjects are attempting to categorize the objects by: either rough details (low spatial frequency) or fine details (high spatial frequency) can be used, depending on what is being judged. For instance, when judging where a rapidly presented image on a screen is a beach or not, it appears that people use large areas of color (‘blobs’) to make rapid, course judgments; in contrast, more attention is paid to other, subtler details when finer-grained judgments are required. Similar results were found by Goldstone & Styvers (2001) when using blended images of faces to vary dimensions. Categorization of faces was improved when images shared either similar or dissimilar features, evidently due to the more careful attention paid to images that shared them. Rosenbaum et al. (2001) as well show indicates strong ties between perceptual/motor skills and cognitive manipulation of the resulting stimuli. Studies such as these speak to the close ties between perception of environmental stimuli and the cognitive processes that become aware of the objects and processes producing those stimuli (see also Pylyshyn 1999 and comments in that issue).

Before discussing features themselves, several background comments concerning words and their use are in order.

3 On word meaning

This section makes several points about words and their use. Words are discussed in terms of their labeling function, their componential nature, their organization, and their effect on perception. In addition, the use of adjectives as reflections of attribute assignment is discussed.

3.1 Words are labels applied to mental constructs

It was noted in §2.1 that much of learning appears to be constructing some manner of mental representation of our world. Taking this one step further, vocabulary acquisition can be viewed as labels being attached to mental constructs of objects, attributes or actions. More specifically, word learning and word use can be seen to be a culturally-mediated application of labels assigned to entities in the world perceived through our sensory input. This can be represented by the following figure:

physical properties of objects;
language input

⇓

<senses & sensory filters>

⇓

perception
(i.e. low-level awareness)

⇓

<filters due to cultural/linguistic
experience>

⇓

linguistic acquisition & use

Figure 1. A simplistic filter model of object feature labeling.

This reflects the reality of investigations such as the current one; while we can measure attributes of physical objects (weight, length, luminescence, etc.) and would like to be able to detail perception of physical properties of objects (e.g. the manifestation of assigning a relatively long item the label *long),* we cannot directly do so; until such a time as we can unravel the neurophysical manifestations of feature perception in the brain’s circuitry, we must rely on indirect observations of brain activity and language use. We may shoot for detailing perception; we must for the present continue to settle for detailing and interpreting linguistic use.

The assignment of labels to concepts appears to be a universal characteristic of language acquisition. Word acquisition occurs quite naturally even in cultures where children are not spoken to directly or where objects are not named by adults (Bloom 2000: 7-9) and with blind children (Landau & Gleitman 1985). What it truly astounding in the latter case is that blind children develop vocabularies very comparable to sighted children of the same age learning the same language, including spatial and even color terms (Landau 2000). In a striking example of blind children’s mastering of spatial concepts, telling a sighted child to look up results in them tilting their head back to gaze upward; telling a blind child to look up results in them reaching up with their hands while leaving their head level (Landau & Gleitman 1985:56); touch is their primary mode of “looking”. Further, both deaf and hearing children, whether being exposed to only one mode of linguistic input or growing up in a mixed spoken/signed environment, acquire all significant linguistic milestones at the same pace (Petitto 1984, 1987, 1988, 2001; Petitto & Marentette 1991, Petitto et al. 2001).

The fact that all children, whether learning words via both modalities simultaneously or via one modality only, do so readily speaks to cognitive principles guiding vocabulary acquisition and use. This use occurs even when the stimuli being monitored through one of the usual primary senses has been circumvented.

Having said that, attention will be turned to the nature of the labels that we apply to the concepts reflecting objects in our environment.

3.2 Word meaning contains a componential feature component

Words are not, of course, simply holistic labels applies to concepts. There seems to be good evidence in favor of maintaining a feature-based component to word meaning.

Katz & Fodor's (1963) seminal paper on the componential structure of lexical entries provided us with what has been termed the 'checklist' theory of word meaning. That is, when labeling an object or concept, a check of the features of both the object/concept is made against the features inherent in words; when a match is found, the word can be applied. Besides lexical processes, this type of feature matching can be applied to other levels in the same way (e.g. segmental feature matching when applying sound rules; lemma feature matching in code switching as in Wei 2000; etc.). In contrast, the work of Rosch (1973, 1975), Rosch et al. (1976), Fillmore (1975) and Coleman & Kay (1981) gave us a 'prototype' theory of word meaning—that word meaning is not a simple checklist describing the necessary and sufficient features that the word represents, but instead word use is governed by category exemplars, or prototypes, that best reflect the word being used or recognized.

Miller & Johnson-Laird (1976: 326-328) maintain that something in the middle is most likely the case. In this sort of hybrid analysis, prototypes would guide us in rough judgments, while features allow fine-grained distinctions to be maintained. This would parallel the use of both fine- and course-grained features when recognizing complex scenes (e.g. Oliva & Schyns 1999; Schyns & Oliva 2001).

New intriguing support for component features of words has been provided by psychological studies. Kellenbach et al (2000) looked at activity in the portion of the brain involved in semantic processing when presenting word pairs that shared semantic features; e.g. 'coin' and 'button'. In pairs such as this the only association between the words is the visual attributes of the referent objects—in this case, both are round. They found that 'button' presented before 'coin' did not make recognizing 'coin' any faster, but it did result in a significantly faster activation of the portion of the brain associated with the interpretation of semantic visual features. This strongly suggests that during lexical recall (i.e. determining what concept or concepts a word is labeling) the shape of a referred object is called up. This lends strong support to the existence of features as a permanent part of our lexical knowledge, whether their manifestation be some sort of checklist or an inherent part of a stored representation of an object.

3.3 Word meaning is a hierarchical

We do not simply maintain a large list of labels (i.e. words) for concepts, of course. Words appear to be arranged in some sort of hierarchical network, with more central (i.e. prototypical) words standing in a central relationship to more peripheral items. As noted in Miller (1998), despite the many criticisms that can be made about hierarchical accounts of lexical orientation, it still seems to be the best descriptive mechanism for nominals. This is why WordNet continues to arrange nominal items in a hierarchical fashion. And as noted in Rosch et al. (1976) provided good evidence for this hierarchical organization as well. They showed that items can be roughly sorted into four basic categories, with other items arranged within those four.

This hierarchical organization is also supported by priming studies, whereby a target item is preceded by a related item during presentation. If the target item is consistently recognized faster when preceded by the similar item (i.e. has been 'primed' by the preceding item), the two items can be said to be closer related semantically. Studies involving priming have indicated hierarchical organization of our mental lexicons. Intriguingly, Rossell et al. (2001) measured semantic priming not just with reaction times but with active measurement of brain activity (functional Magnetic Resonance Imaging, or fMRI) to more closely indicate areas involved in semantic priming; these areas can now be targeted in further study.

3.4 Word learning affects mental organization

Word learning is, of course, not as simple as assigning labels to concepts that we encounter in the environment. As will be discussed further below in §4.7, our language ability not only reflects our growing linguistic ability, it also appears to causally affect our ability to make judgments and interpret our environment.

Landau (1984) notes that learning a language (e.g. Korean) will cause the child to focus on spatial relationships and object properties that are highlighted by that language's vocabulary; these relationships and properties will be different than a child learning a different language (e.g. English). Bowerman (1994), in her discussion of inter-language differences in rendering spatial terms, notes the same phenomenon. Different language's terms for different spatial relationships are quite different, each selecting from a different subset of possible features for the definition of a specifically-labeled relationship (e.g. *on* in English relies on contact without reference to an object's shape or spatial relationship with what it is contacting; both German and Dutch break these relationships up into 3 different prepositions, each with different relational distributions). Speech errors in research show that while the relationships attended to might reflect basic but differently selected cognitive principles of organization, they must be salient relationships or the child will not apply them in their language use. The same point was made in L. Bloom (1993), termed the *principle of relevance:* "Words are learned when they are relevant to what the child has in mind." What the child has in mind, of course, is influenced not only by their immediate environment, but also what is being presented by and attended to by their caregivers—all of these which are influenced by the culture they are growing up in.

In studies on how both children and adults orient themselves in a dimensional space, Hermer-Vazquez et al. (1999, 2001) and references therein have indicated that the ability of older children and adults to reorient themselves in a room for the task of finding objects depends on, as much as anything, the ability to talk about the dimensions involved in the space. In addition, Hsieh et al. (2001) have demonstrated that different areas of the brain are used to process speech sounds for Chinese (Mandarin) and English speakers; their findings indicate linguistically-mediated differences in how the languages are parsed prosodically by speakers.

Note that this is not the same as support for what is commonly called the Sapir-Whorf hypothesis. It is not saying that learning a certain language will make one more able or unable to discern a given attribute; it is saying simply that learning a given language will simply make certain properties and attributes more salient since they are being focused on by the surrounding speech community.

Discussion will now turn to the use of adjectives in the current line of research.

3.5 Adjectives as feature assignors

One of the stumbling blocks run into in previous analyses of componential features is the sheer complexity of the data. This is eloquently demonstrated by decomposition of the simple assertion “Mike is tall” adapted from Miller & Johnson-Laird (1976: 326), where ‘Mike’ is x:

(1) HUMAN (x) & MALE(x) & GREATER (VERTICAL (DISTANCE FROM (earth, TOP(x))), NORMAL (VERTICAL (DISTANCE (earth, TOP (HUMAN MALE)))))

In other words, “x is a human male, and the vertical distance from the earth to the top of x is greater than the normal vertical distance from the earth to the top of a man”.

The time-honored way to limit the complexity of a phenomenon is, of course, to simply restrict one’s view to a subset of the data. Such is the case in the present line of investigation.

As noted in Dixon (1982), the vast majority of the worlds languages have, in addition to the evidently mandatory categories of nominals and verbals, a class of adjectivals. To further restrict the attribute categories of Dixon (1982: 15ff)[3], it makes sense to start with adjectivals as representative of the attributes that can be assigned to a given concrete object. The current research therefore focuses on features that can be assigned to a concrete physical object, supplemented by spatial and temporal terms—spatial organization appears to be an intrinsic human cognitive trait (Landau 1994), and temporal sensitivity to our diurnal cycles appears to be present in all life forms exposed to the nature order of night and day (Young 2000).

As a starting point, the cogPrime database contains semantic features assigned to concrete objects and their attributes, supplemented by a set of meta features for use in organizing the data set, and features reflecting an object’s position in time and space.

---

[3] Dixon 1982 included the category value for adjectives such as *good* and *bad*, as well as the category HUMAN PROPENSITY for adjectives such as *jealous*, *loyal* and *merry*.

Attention will now turn toward the main focus of the paper: features.

4 Features

This section discusses the use of the term ‘feature’ in several contexts, beginning with its long-established use as a property of a phonological segment.

4.1 The use of features

The word *feature* can be used in many senses, and is present in virtually every linguistic field. Features have a particularly long and strong history in the fields of phonology and phonetics (Trubetskoy 1939/1969; Jakobson et al. 1952; Chomsky & Halle 1968) where features are generally associated with both articulatory movements and aural impressions (Keating 1988). Indeed, a large part of the evolution of Generative Phonology (e.g. Lexical and Prosodic Phonology; see Hargus & Kaisse 1993 and Nespor & Vogel 1986) has been due to detailing the application of phonological rules to segmental features and the domain of rule application. Investigations into segmental features has led to the emergence of the hierarchical organization of features in Feature Geometry (Clements 1985; McCarthy 1988; Sagey 1986, 1988).

Other fields in addition to phonology and semantics (see §3.4) have made good use of features as well. Standing in opposition to segmental analyses such as those mentioned above, at the level of discourse analysis Wei (2001) has shown that code switching, whereby bilingual Chinese/English speakers insert in tact phrases from one language into a matrix sentence of the other, occurs only when the syntactic, semantic and morphological features of both the sentence and the inserted phrase agree. In this case *feature* is used as all of the informational items that must be checked in the overall framework (the lemma) of the utterance.

Features have also been used successfully to facilitate delineation of word sense relations. By reference to word meaning as sets of features, synonymy, hyponymy and hypernymy can be readily explained: two lexical items that contain or reference the same set of features are synonyms; a hypernym contains or references fewer features (i.e. is more generic) than one of its related hyponyms (see, e.g., Katz & Fodor 1963 or Lyons 1968 on extension vs. inclusion). As a concrete example, the difference between *stamp* and *stomp* in the author’s dialect of American English (and assumedly, the broader English-speaking community) can be detailed as follows:

(2a) *stamp:* to transfer energy to the floor or ground by virtue of kinetic energy due to motion of the foot

(2b) *stomp:* to transfer energy to the floor or ground by virtue of momentum due to inherent weight

This is exemplified by the following examples:

(3)	Quit stamping around!	Quit stomping around!
(4)	Quit stamping your foot!	*Quit stomping your foot!

(5) a small child stamping/*stomping around the room

(6) a boy/girl/man/woman stamping/stomping around the room

This difference could be attributed to a difference in the features associated with each label:

(7a) stamp: energy imparted due to [+motion]

(7b) stomp: energy imparted due to [+inherent weight]

Features such as these allow us to directly detail such differences.

Psychological studies in particular have investigated the nature of features (in its more traditional use of a physical trait) and their use in categorization (see also §2 above). They have been shown to guide facial recognition (Young et al. 1997; Schyns & Oliva 1999), including whether someone is angry or not[4] (Lundqvist 1999) and whether someone is from another race (Levin 2000). Prosodic cues and pauses have been shown to be good measures of clinical depression (Alpert et al. 2000); this is entirely in line with normal experience where we routinely (even if subconsciously) monitor the prosodic characteristics of those we talk to to determine their state of mind, be it depression, anger, sadness or mistrust. In addition, Hermer-Vazquez (1999, 2001) and references therein detail the use that children, adults, and even rats make of

---

[4] Surprisingly, the eyebrows seem to provide the greatest cue to anger, the other facial features being secondary or modulating features.

environmental landmarks for orienting themselves in their environments.

One thing runs central in all frameworks that make reference to features, even if it is only implied or simply ignored: features are collections of experiences or impressions somehow encoded in the brain (i.e. some sort of memory), whether it be subconscious parsing of the speech stream or attending to details of a complex picture. Features are physically neural activity triggered by environmental stimuli. The existence of features can be inferred by responses to presentation of stimuli (words, pictures, etc.). We cannot yet directly observe this neural activity patterning. However, Kohonen & Hari (1999) have provided intriguing models of how features might be organized and observed in the brain. Pending results of the methodology they propose, we remain forced to resort to more classical methods for determining their existence and use.

Attention will now be turned to just what the mental organization of such features might entail.

4.2 Mental organization of features

As noted above in §3.1, we cannot yet directly observe the manifestation of features in the brain's circuitry. However, by creative use of language paradigms, indications can be gleamed.

Such creative use can be seen in Solomon & Barsalou (2001). They tested participants with similar-category items, things labeled with the same lexical item which are fairly distinct (e.g. *mane* and *belly* for various animals). For example, the mane of a horse and the mane of a male lion are not particularly similar at all. Their prediction was that if a single, global representation of a word such as mane existed, then either verifying that a horse has a mane or a lion has a mane should provide the same advantage when later verifying that a pony has a mane. However, if representations of words are local—that is, a horse's (and pony's) mane involves a distinct representation separate from that of a lion's mane—only verifying that a horse has a mane should help with the later verification of the pony's mane. Identifying that a horse has a mane should not help in later identification that a lion does, since at the local (i.e. fine-featured) level the two types of manes are quite distinct. Other lexical items with physically distinct manifestations were used as well.

Their findings indicated that such representations are indeed local; the mane of a pony and a horse is not closely associated with the mane of a lion. This held true across lexical items presented (i.e. the handle of a suitcase, hammer and car door). They concluded that recognition and processing of properties is complex—as one would expect—but nevertheless is grounded in mental representations of the features being identified: representations of parts of objects such as the mane of a horse are stored with physical detail found in the real-world referent, and these representations are actively searched through when identifying those referents. In simplistic terms, features are stored locally (i.e. as a part of what they describe) at least when necessary or felicitous, not simply globally (i.e. as a holistic feature subsuming all instances).

4.3 Underspecification of semantic features?

As discussed in §3.5 above, one way to facilitate analysis of a complex phenomenon is to limit the scope of the investigation—in this case, restricting one's focus to features modifying physical properties of concrete objects. Another way to limit the complexity of such formulations would be to appeal to default readings for various categories being referenced. For example, in reference to the example 'Mike is tall' in §3.5, in absence of other evidence an object being referenced by a proper noun is a human, and references to height always involve the ground plane so that it not need be specified—underspecification of semantic features, if you will. Since our cognition appears to involve some sort of representation of the real world, there is no reason to assume that the assertion "Mike is tall" does not simply call forth a generic image of a human male of average height.[5] The object reflecting Mike is compared to this generic man, and seen to be taller.

This calling forth of a 'generic object' is entirely in line with the results of prototype semantic theory (Rosch 1973, 1975; Fillmore 1975; Rosch et al. 1976; Coleman & Kay 1981)—where a concept such as 'bird' is represented not by a list of features ([+has wings], [+flies], etc.) but by prototypes or focal instances of the concept (e.g. a robin exemplifies *bird* much better than *penguin* does). This parallels current programming technique, object-oriented programming, which makes heavy use of *object code*.

[5] 'Average' being, of course, a function of one's experience and situation—I am very tall in most of Asia, of only of average height in America or Europe.

Instead of relying on the subroutines of yesteryear to achieve programming tasks, the programmer defines objects and classes of objects. These objects can be primitive types, carrying very little information with themselves, or quite complex, containing within themselves information from a great number of other objects. As a concrete example, a class defining a square would specify only that all sides and angles are equal; further information would be provided by referencing the class of *quadrilaterals* which would be specified as having four sides, and so on. In addition, most programming languages contain methods for modifying a generic object with specific features so that it differs from the prototypical instance of the object. Utilization of such programming techniques would seem to hold great promise for future modeling of attribute assignment, and will be pursued in further research.

4.4 The permanence of features

In addition to the local nature of feature storage noted above, the permanence of features is suggested by the persistence of objects in our subjective reality. While conventional wisdom and most researchers hold that this permanence of mental constructs is indicative of discrimination of boundaries, colors and object surface patterns (e.g. Spelke 1982, 1983, 1994; Kellman & Spelke 1983; Spelke & Van de Walle 1993; Hermer-Vazquez et al. 1999), others argue for the permanence of objects in our mental representations being due to relative motion or the lack thereof—moving objects produce patterns that tell us they are moving; stationary objects produce patterns that tell us nothing is moving (Shipley & Kellman 1997; Shipley 2000). This proposal is entirely in line with both those and previously mentioned studies, and provides an intriguing avenue of further research into object permanence as well as motion, especially in light of the fact that detecting boundaries of objects (alternatively, fronts of movement) appear to be hard-wired into the primate brain (e.g. Janssen et al. 2000). It remains to be seen whether this sort of pattern identification can be adapted to the tactile mode necessary to allow it to be used when blind children are learning language (see the discussion in §3.1 above).

Whichever the case may be, permanence of features is indicated, either as assigned attributes of the objects we construct, or as attributes of the patterns that indicate object permanence. Permanence of features is also supported by a wealth of anecdotes relating to children's language acquisition. For example, while learning to talk A (app. 2;3) asked her father to give her the toy pig on the shelf by reaching and calling for it. Because the toy in question had a very cat-like face, and was turned face toward them, her father did not understand her request. After becoming frustrated and (in effect) scolding him on his stupidity, A fetched the toy herself, turned it around and indignantly showed her father its curly tail.

Somehow when learning her first word for pig some sort of feature such as [+curly tail] became salient and was given prominent status, to the point that very cat-like ears, eyes and nose were no longer pertinent—their presence was not sufficient to override the fact that cats never have spiral tails, but pigs always do (at least in her then subjective world). This is despite the seemingly central role that recognizing facial features plays in human recognition (see §3.1 above). The toy was a pig, period.

Anecdotes such as attest to the child attuning to a systematic analysis of the speech stream and surrounding environment they encounter. Categorizing objects based on features of those objects would seem to be a part of language acquisition strongly relied on by the child.

4.5 Features as a tool for language acquisition

The categorization of objects discussed in §2 above serves a very important function in language development: it allows even very young children to infer facts about their environment. This is supported by current research. Hespos & Baillargeon (2001) showed that children as young as 2 mos. are sensitive to whether or not something can be put into a container; Needham (1999) showed that children as young as 4.5 mos. attend to the shape but not so much texture of objects in determining whether to group objects together as they move.

Goldstone et al. (2001) indicates that categories and features are a bi-directional system for categorizing things in the environment; indeed, they present evidence that not only are category ‘prototypes’ are important for categorizing objects, but that the relevant features used to compare objects within a given category are emphasized. Just as a cartoonist highlights certain features to make a celebrity immediately recognizable from a sparse representation, people appear to make use of the same sort of mechanism to facilitate learning a novel category. If adults make use of this highlighting process, it is quite likely that the same mechanism is available to the young language learner, and would greatly facilitate making real-world associations.

4.6 Features as reflective of scales

Adjectives as reflections of attribute scales has a long history in the literature as well (see, for example, discussion in Weinrich 1972; Dixon 1982; Lyons 1968, 1977; Jackendoff 1983). Cruse (1986) in particular makes clear the relation between scale endpoints and the 'pivotal' or neutral region in the middle of the scale; for most adjectives in English there are not lexical items associated with the middle of the scale.[6] Although one could conceivably find an object whose length is neither short nor long, there are no lexical terms for 'average length' in most languages. For studied languages, the same holds true for other lexical items (i.e. adjectivals) reflecting object attributes.

Features such as those reflected by such adjectivals can be represented by the use of axes. For example, 'long' in English refers to the principle dimension of object, regardless of its orientation in space.[7] For most scales and axes, there appear to be no static regions (Lyons 1977: 273-275), but the use of these axes, whether they are physically delineated or not, holds even with young children (e.g. Landau 1994, 2000).

Finally, we routinely and actively adjust the scales by which we measure and interpret our environment according to the task at hand. This occurs with low-level sensory perception (e.g. the judging of a color as red whether in bright sunlight or dim candle-light), and with higher-order processing (e.g. judging a child big when standing next to a younger classmate, but small when next to an elephant).[8]

4.7 Features as 'living' entities

Another characteristic of features that has become apparent is that while their use may be in inherent part of our intellect, the set of features used and the nature of each feature in that set is not fixed. Features are fluid; they are living entities that respond to our increased understanding of our environment. This has been demonstrated

---

[6] Temperature modifiers are one exception; many languages do divide the temperature scale up into several regions; e.g. English *cold, cool, warm, hot.*

[7] The exception, of course, is when the principal dimension corresponds to what we refer to as the height of something.

[8] See Chater & Brown 1999 for discussion of scale invariance—in simplistic terms, the ability to use the same scale in any situation by simply adjusting the values of the endpoints of the scale—as a "unifying psychological principle" from which several physical laws derive.

thoroughly in the realm of auditory discrimination. While categorization of aural input appears to be an innate feature of the animal kingdom (Kuhl & Miller 1978; Kuhl 1987; Wyttenbach et al. 1996; Doupe & Kuhl 1999), the categorization of speech sounds is highly influenced by our linguistic experience (e.g. Werker & Tees 1984; Werker 1986, 1994; Werker & Pegg 1992; Best 1994).

This 'training' seems to begin quite early, most likely during the prenatal experience see the papers in Mehler 1983; Mehler & Fox 1985; see also Locke 1993:34 *ff.).* Day-old infants show sensitivity to their mother's voice, actively sucking faster or slower as needed to hear a recording of their mother's voice instead of another woman's, and then less than a day later in a second trial reversing the sucking rate to again hear their mother's voice (DeCasper & Fifer 1980). This is despite being housed in a group nursery and being exposed to other women's voices. Prenatal sensitivity only makes sense, since the unborn child's ears are fully functional several weeks before birth. Experience hearing even their mother's and acquaintances distorted voices give the child an opportunity to begin categorizing their environment in a systematic way as early as the age of 1 month (Eimas et al. 1971; Eimas 1974, 1975, 1985; see Eimas et al. 1987 for a review).

Fluidity of features has also been demonstrated in a large number of cognitive studies. Lesgold et al. (1988) demonstrated learning of novel features when participants were trained to read x-rays; Schyns & Rodet (1997) used circles containing DNA-like squiggles, termed "Martian cells", to show that people can likewise learn to categorize novel objects using learned features.

At the category level, Ross & Murphy (1999) showed that different categories can be used to judge the same food items at the same time (e.g. a pizza can be simultaneously categorized as junk food or Italian food), suggesting fluid category use. This fluidity extends to the scales used to judge features: as discussed above, Schyns & Oliva (1999) used composite facial images (images of one face superimposed over another) to show the scale used to judge an image flexibly adjusts to the feature being judged.

Rehder & Ross (2001) also supports the notion that the categories objects are grouped into are themselves fluid. Their work finds that human cognition is filled with abstract categories that people use to classify things. These abstract classifications are necessarily the result of life experience, not some innate ability—there is no universal need for the types of classifications people can make with regard to social groups, societal institutions, and legal, political, and military activity. Matan & Carey (2001) showed that 6 year olds pattern with adults in judging the function of an object (e.g., something that could be either a teapot or a watering can) based on weighing what they were told was its original function with what they were told it was being used for currently, whereas 4-year old based their categorizations on the order the choices were presented; i.e. they judged something to be a watering can if that was the first intended use of the item they heard, or as a teapot if they heard that use first, regardless of what it was currently being used for. The study conducted by Nazzi & Gopnik (2001) suggests by 20 months words are used to begin forming categories. Hermer-Vazquez et al. (1999, 2001) showed that children's ability to reorient themselves in a fixed-shape but differential-feature environment (e.g. a room with one wall colored) kicks in around the age of 5-7, after which they display fairly adult ability to both reorient themselves and find a hidden object in the room. Finally, Goldstone et al. (2001; see discussion in §3.4 above) provided evidence of changes in both the categories into which objects are assigned and mental representations of objects based on highlighted physical attributes. They found that changes in category labels are supported when making judgments about objects belonging to different categories (the *Strategic Judgment Bias* account) and changes in the features of the objects themselves when making similarity judgments about objects in the same category (the *Changed Object Description* account).

As for whether this ability to define new categories is species-specific or not, Fagot et al. (1998) studied attention shifts in both baboons and humans when encountering new items (in this case, hieroglyphic-like symbols). Their findings indicate differences in how much attention each species gives to visual cues. Their best-fit model suggests humans retain the features previously encountered, and instead shift their attention to novel features appearing in the stimuli. Baboons did not show such shifts, but instead re-processed each appearing symbol. This human ability to rapidly focus attention on novel features would seem to be critical for the skill of categorizing novel stimuli.

All of these indicate that features we use to categorize objects, and the categories we group them into, actively interact to produce our version of our environment.

5 Representing features

One last topic related to features will be discussed in this paper, since it is germane to the database of semantic features being constructed—the manner in which the mental attributes being discussed can be represented.

5.1 Natural language labels

To investigate the use of features as lexicon organizers and filters for input, a multi-lingual database of features needs to have a way to organize the data itself. Such a database must have, of course, labels for the features contained therein. As noted by Lyons (1968: 479-480), using natural language to talk about features is a risky proposition—the substantial bias of one's native tongue and personal experience make judgments suspect.

The goal, of course, is to delineate a list of features used by humans to refer to objects in their environment. The reality is that we end up with a list of words and/or phrases that we assign as labels for those features, all of which reflect the language they are written in. In other words, it is not possible to assign a feature the label [+height] without calling forth one's personal and cultural experience with what is normally considered tall, and what range of height values exceptions might take. Scientific terminology can help reduce this bias somewhat, at the expense of making feature descriptions difficult to understand for both native and non-native speakers at best and impenetrable at worse.

To avoid this cultural/ideolectal prejudice, a different tack seems in order.

5.2 Alternate representations for features

In this day and age the vast majority of researchers in any field make active use of the internet, either in their work, their private life, or both. We are all familiar with the multitude of images that barrage us on-line. While many of these images contain text that help us achieve some goal (e.g. navigation or search buttons), or are intended to influence our on-line behavior (e.g. banners inducing us to make a purchase or access another site), we also can find images that are merely decoration—the design part of a web page. Such images not only can be appealing, but make use of web pages easier since the average human is so highly visually dependent (see Cruz & Leveille 2001 for a natural extension of this use of images, an image-based search engine; see also Karp

2001 for making sense of the otherwise impenetrable multitude of gene functions associated with E. coli via the use of interlinked images).

Since the current project contains a web-accessible component, it seems natural to make good use of images in the elicitation of data. The use of images has a long history in image-recognition and processing studies (e.g. Oliva & Schyns 1997, 2000; Schyns & Oliva 1999) and has been used in the past for phonological/phonetic studies (e.g. Davis 1994, 1995). Surprisingly, this seems to be a little-used means of eliciting semantic features (cf. the text-based elicitation of features used in Ross & Murphy 1999; Hinojosa et al. 2001; Garrard 2001).

Perhaps the reason is that any image carries with it an incredible amount of detail. For example, an image of a face can be categorized in many different ways (Schyns & Oliva 1998). With objects used to represent basic attributes, the problem is magnified; an image of a bowl of water might be interpreted as a stimulus for either its size, the volume it holds, or whether or not it is full or empty. To avoid the confounding of stimuli, line drawings of the type used in Damian et al. (2001) are thought to be better; by limiting the image itself to rough outlines representing the attribute in question, a more limited categorization should be able to be achieved. This is in line with the highlighted cognitive representations of Goldstone et al. (2001), discussed in §4.5. Line drawings therefore form the foundation of the elicitation materials being developed.

However, there are several potential problems with using line drawings for measure of such things as width. First, one must be certain that the participant is identifying the correct dimension of the object in questions—for *short/long,* the primary (usually longest) axis, for *narrow/wide*, the secondary axis. This possibility of the participant rotating the image in their mind must be circumvented. To prevent such rotation, vectors can be added to the drawings to assist in identifying the axis in question.

Secondly, it is difficult to provide a person with a line drawing that is unequivocally long or wide; will 3 meters suffice, or will it take 10 or 30 meters width to ensure that all participants are identifying the desired end of the scale in question? A 10-meter drawing would be both difficult to produce and difficult to present. To help ensure correct identification of the desired region of an attribute scale, the two figures can be juxtaposed; limited, simplistic natural language phrases can be used to unambiguously identify the object in question, and units of measure can likewise be included to establish the objects in relative frameworks rather than literal ones. By careful limitation of both the figures and any language that accompanies them, it should be possible to ensure elicitation of the desired attributes.

As an example, the following stimuli could be used to elicit the equivalent of the English pairs *short, long* and *narrow, wide:*

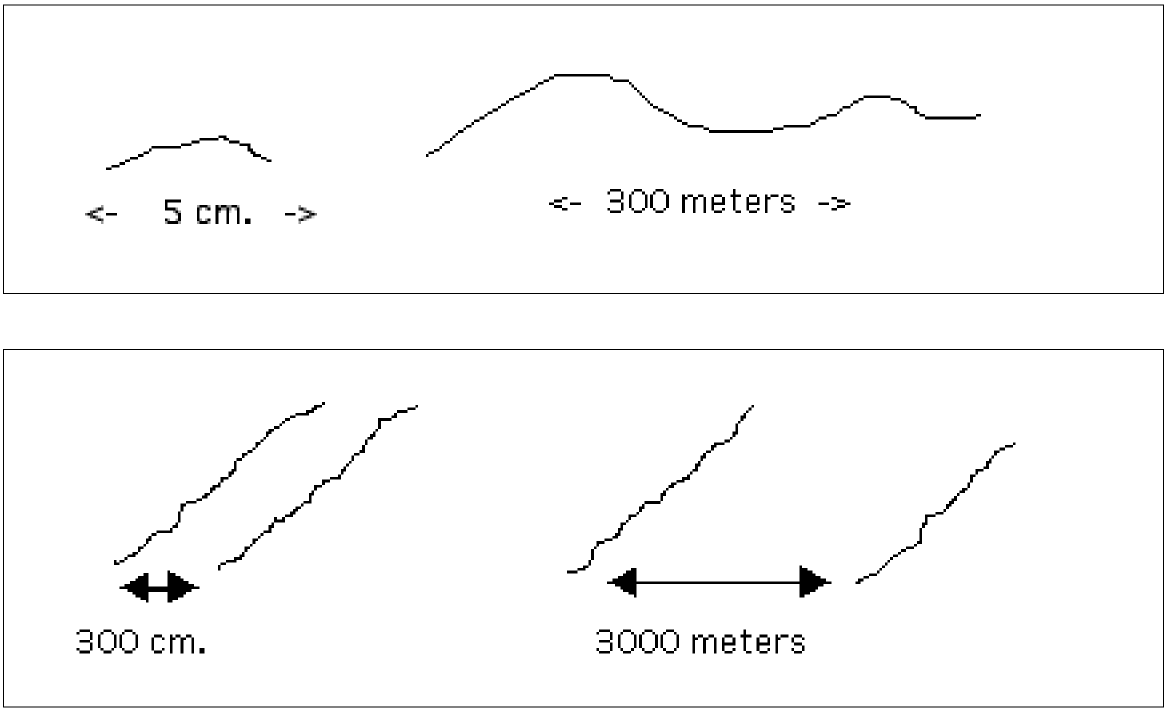

Figure 2. Elicitation materials for *short, long, narrow,* and *wide.*

The included vectors should help the participant unambiguously identify the dimension in question, while the units of measurement used in the figures should help participants establish the relative size of the object being labeled. An alternative to the units of measurement would be a reference object, whose size varied according to scale.

For elicitation of color terms especially, the use of images is called for. Instead of simply asking a participant to translate *dim* and *bright,* or show them a picture of a lamp, a gradated series of images makes the contrast quite clear. For example, the follow image could be used for the feature category BRIGHTNESS:

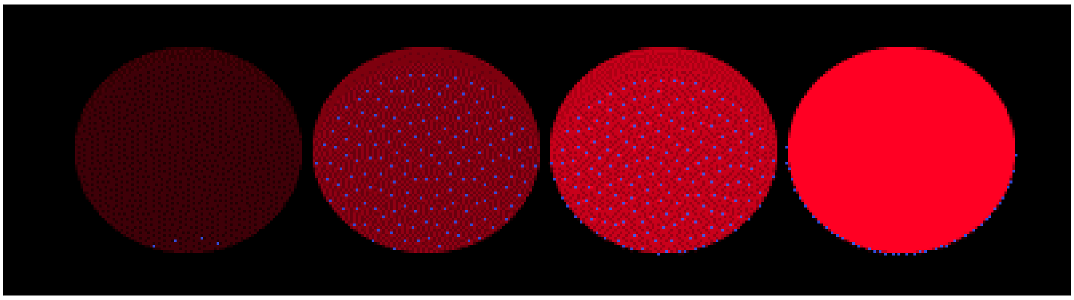

Figure 3. Elicitation image for *dim* and *bright.*

Once images have been created for all desired features, they can be embedded in a web page for direct solicitation of data via the web. This will allow those interested (and/or paid) to participate at their own pace, according to their own schedule. Such web-based data collection may prove to be useful in other areas of investigation as well.

6 Summary

Current research on features and categorization speaks to a higher-order process being involved that likely spans disciplines and sensory modalities. Concentrating this higher-order processing may continue to provide insights into just how much overlap is involved. In particular, a more vigorous pursuit of cross-over between the processes indicated by phonological and phonetic features processing (e.g. underspecification; Feature Geometry) and the features used to categorize other sensory modes (e.g. visual and tactile sensations) would seem to greatly benefit all fields of investigation. Finally, it was suggested that using images as an alternative to natural language for the elicitation of the semantic features bears further consideration.

7 References

Ahn, Woo-kyoung, Nancy S. Kim, Mary E. Lassaline and Martin J. Dennis. 2000. Causal Status as a Determinant of Feature Centrality. *Cognitive Psychology* 41: 361–416.

Alpert, Murray, Enrique R. Pouget and Raul R. Silva. 2001. Reflections of depression in acoustic measures of the patient’s speech. *Journal of Affective Disorders* 66: 59–69.

Best, Catherine T. 1994. The emergence of native-language phonological influences in infants: A perceptual assimilation model. In eds. Judith C. Goodman and Howard C. Nusbaum, *The Development of Speech Perception: The transition from speech sounds to spoken words*, 167-224. Cambridge: MIT Press.

Bloom, Lois. 1993. *The Transition from Infancy to Language: Acquiring the Power of Expression*. New York: Cambridge University Press.

Bloom, Paul. 2000. *How Children Learn the Meanings of Words*. Cambridge: The MIT Press.

Borden, Gloria J., Katherine S. Harris and Lawrence J. Rapheal. 1994. *Speech Science Primer: Physiology, Acoustics, and Perception of Speech*. Baltimore, Maryland: Williams & Wilkins.

Bowerman, Melissa. 1994. Learning a semantic system: What role do cognitive predispositions play? In ed. Paul Bloom, *Language Acquisition: Core Readings*, 329-363. Cambridge: MIT Press.

Chater, Nick & Gordon D.A. Brown. 1999. Scale-invariance as a unifying psychological principle. *Cognition* 69: B17–B24.

Chomsky, Noam and Morris Halle. 1968. *The Sound Pattern of English.* New York: Harper and Row.

Clements, George. 1985. The Geometry of phonological features. *Phonology Yearbook* 2: 225-252.

Coleman, Linda and Paul Kay. 1981. Prototype semantics. *Language* 57: 26-44.

Cruse, D. A. 1986. *Lexical Semantics*. Cambridge: Cambridge University Press.

Cruz, Isabel F. and Peter S. Leveille. 2001. As You Like It: Personalized Database Visualization Using a Visual Language. *Journal of Visual Languages and Computing* 12: 525-549.

Damian, Markus F., Gabriella Vigliocco and Willem J.M. Levelt. 2001. Effects of semantic context in the naming of pictures and words. *Cognition* 81: B77-B86.

Davis, Katherine. 1994. Stop voicing in Hindi. *Journal of Phonetics* 22: 177-193.

Davis, Katharine. 1995. Phonetic and phonological contrasts in the acquisition of voicing: voice onset time production in Hindi and English. *Journal of Child Language* 22: 275-.

DeCasper, Anthony J. and William P. Fifer. 1980. On human bonding: Newborns prefer their mother's voices. *Science* 208: 1174-1176.

Dixon, Robert M. W. 1982. Where have all the Adjectives Gone? and Other Essays in Syntax and Semantics. The Hague: Mouton.

Doupe, Allison J. and Patricia K. Kuhl. 1999. Birdsong and human speech: Common Themes and Mechanisms. *Annual Reviews: Neuroscience* 22: 567-631.

Eimas, P. D., E. R. Siqueland, P. Jusczyk and J. Vigorito. 1971. Speech perception in infants. *Science* 171: 303-306.

Eimas, Peter D., Joanne L. Miller and Peter W. Jusczyk. 1987. On infant speech perception and the acquisition of language. In ed. Stevan Harnad, *Categorical Perception: The groundwork of cognition*, 161-195. Cambridge: Cambridge University Press.

Eimas, P. D. 1974. Auditory and linguistic processing of cues for place of articulation by infants. *Perception & Psychophysics* 16: 513-521.

Eimas, P. D. 1975. Auditory and phonetic coding of the cues for speech: Discrimination of the [r-l] distinction by young infants. *Perception & Psychophysics* 18: 341-347.

Eimas, Peter. 1985. The perception of speech in early infancy. *Scientific American* 252: 19-23.

Etcoff, N.L., Maggee, J.J., 1992. Categorical perception of facial expressions. Cognition 44, 227–240.

Fagot, Joël, John K. Kruschke, Delphine Dépy & Jacques Vauclair. 1998. Associative learning in baboons (Papio papio) and humans (Homo sapiens): species differences in learned attention to visual features. *Animal Cognition* 1: 123–133.

Fellbaum, Christiane ed. 1998. *WordNet: An Electronic Lexical Database*. Cambridge: MIT Press.

Fillmore, Charles. 1975. An alternative to checklist theories of semantic meaning. In *Papers from the First Meeting of the Berkeley Linguistics Society,* 123-131. Berkeley, Berkeley Linguistics Society.

Garrard, Peter, Matthew A. Lambon Ralph, John R. Hodges and Karalyn Patterson. 2001. Prototypicality, distinctiveness, and intercorrelation: Analyses of the semantic attributes of living and nonliving concepts. *Cognitive Neuropsychology* 18: 125-174.

Gokcen, Jeanne M. and Robert A. Fox. 2001. Neurological evidence in support of a specialized phonetic processing module. *Brain and Language* 78, 241–253.
Goldstone, Robert L. & Lawrence W. Barsalou. 1998. Reuniting perception and conception. *Cognition* 65: 231–262.
Goldstone, R. L. and M. Styvers. 2001. The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology. General* 130: 116-139.
Goldstone, Robert L., Yvonne Lippa & Richard M. Shiffrin. 2001. Altering object representations through category learning. *Cognition* 78: 27-43.
Hargus, Sharon and Ellen Kaisse eds. 1993. *Studies in Lexical Phonology*. New York: Academic Press.
Harnad, Stevan. 1996. The origin of words: A psychophysical hypothesis. In B. Velichkovsky and D. Rumbaugh eds. *Communicating Meaning: Evolution and Development of Language.* NJ: Erlbaum. 27-44.
Hermer-Vazquez, Linda, Elizabeth S. Spelke and Alla S. Katsnelson. 1999. Sources of Flexibility in Human Cognition: Dual-Task Studies of Space and Language. *Cognitive Psychology* 39: 3–36.
Hermer-Vazquez, Linda, Anne Moffet and Paul Munkholm. 2001. Language, space, and the development of cognitive flexibility in humans: the case of two spatial memory tasks. *Cognition* 79: 263-299.
Hespos, Susan J. & Reneé Baillargeon. 2001. Reasoning about containment events in very young infants. *Cognition* 78: 207-245.
Hinojosa, Jose A., Manuel Martın-Loeches, Francisco Munoz, Pilar Casado, Carlos Fernandez-Frıas and Miguel A. Pozo. 2001. Electrophysiological evidence of a semantic system commonly accessed by animals and tools categories. *Cognitive Brain Research* 12: 321–328.
Hsieh, Li, Jack Gandour, Donald Wong and Gary D. Hutchins. 2001. Functional Heterogeneity of Inferior Frontal Gyrus Is Shaped by Linguistic Experience. *Brain and Language* 76: 227–252.
Ikegami, Yoshihiko. 1970. *The Semological Structure of the English Verbs of Motion: A Stratificational Approach.* Tokyo: Sanseido.
Jackendoff, Ray. 1983. *Semantics and Cognition*. Cambridge: The MIT Press.
Jakobson, Roman, Gunnar M. Fant and Morris Halle. 1952 (reprinted by MIT Press 1963). *Preliminaries to Speech Analysis: The distinctive features and their correlates.* MIT Acoustics Laboratories Technical Report 13.
Janssen, Peter, Rufin Vogels and Guy A. Orban. 2000. Selectivity for 3D shape that reveals distinct areas within Macaque inferior temporal cortex. *Science* 288: 2054-2056.
Jusczyk, Peter W. 1992. Developing phonological categories from the speech signal. In eds. Charles A. Ferguson, Lise Menn and Carol Stoel-Gammon, *Phonological Development: Models, Research, Implications*, 17-64. Timonium, Maryland: York Press.
Jusczyk, Peter W. 1994. Infant speech perception and the development of the mental lexicon. In ed. Judith C. Goodman & Howard C. Nusbaum, *The Development of Speech Perception: The Transition From Speech Sounds to Spoken Words*, 227-270. Cambridge: MIT Press.
Kalish, Michael L. & John K. Kruschke. 2000. The role of attention shifts in the categorization of continuous dimensioned stimuli. *Psychological Research* 64.2: 105-116.
Karp, Peter D. 2001. Pathway Databases: A Case Study in Computational Symbolic Theories. *Science* 293: 2040-2044.
Katz, Jerrold J. and Jerry A. Fodor. 1963. The structure of a semantic theory. *Language* 39: 170-210. Reprinted in Fodor, Jerry A. and Jerrold J. Katz eds. 1964. *The Structure of Language: Readings in the philosophy of languages.* Englewood Cliffs, NJ: Prentice-Hall, Inc. pp. 479-518.
Keating, Patricia.1988. A survey of phonological features. Unpublished ms., UCLA.

Kellenbach, Marion L., Albertus A. Wijers and Gijsbertus Mulder. 2000. Visual semantic features are activated during the processing of concrete words: event-related potential evidence for perceptual semantic priming. *Cognitive Brain Research* 10: 67–75.

Kellman, P.J. and E. Spelke. 1983. Perception of partly occluded objects in infancy. *Cognitive Psychology* 15: 483-524.

Kohonen, Teuvo and Riitta Hari. 1999. Where the abstract feature maps of the brain might come from. *Trends in Neuroscience* 22: 135–139.

Kotchoubey, B., E. Wascher and R. Verleger. 1997. Shifting attention between global features and small details: an event-related potential study. *Biological Psychology* 46: 25-50.

Kruschke, John K. 1992. ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review* 99: 22-44.

Kruschke, John K. 1996. Base rates in category learning. *Journal of Experimental Psychology: Learning, Memory and Cognition* 22: 3-26.

Kruschke, J. K. and M. K. Johansen. 1999. A model of probabilistic category learning. *Journal of Experimental Psychology. Learning, Memory, and Cognition* 25: 1083-1119.

Kuhl, Patricia K. and James D. Miller. 1978. Speech perception by the chinchilla: Identification functions for synthetic VOT stimuli. *The Journal of the Acoustic Society of America* 63: 905-917.

Kuhl, Patricia K. 1987. The special-mechanisms debate in speech research: Categorization tests on animals and infants. In ed. Stevan Harnad, *Categorical Perception: The groundwork of cognition*, 355-386. Cambridge: Cambridge University Press.

Landau, Barbara and Lila Gleitman. 1985. *Language and experience: Evidence from the blind child*. Cambridge, MA: Harvard University Press.

Landau, Barbara. 1991. Spatial representations of objects in the blind child. *Cognition* 38: 145-178.

Landau, Barbara. 1994. Where's what and what's where: The language of objects in space. In eds. Lila Gleitman and Barbara Landau, *The Acquisition of the lexicon*, 259-296. Cambridge: MIT Press.

Landau, Barbara. 2000. Language and space. In eds. Barbara Landau, John Sabini, John Jonides and Elissa L. Newport, *Perception, Cognition, and Language: Essays in honor of Henry and Lila Gleitman*, 209-230. Cambridge: MIT Press.

Lesgold, A. M., Glaser, R., Rubinson, H., Klopfer, D., Feltovitch, P. & Wang, Y. 1988. Expertise in a complex skill: Diagnosing X-ray pictures. In M. T. H. Chi, R. Glaser, M. J. Farr Eds., The nature of expertise. Hilldsale, NJ: Erlbaum.

Levin, Daniel T. 2000. Race as a Visual Feature: Using Visual Search and Perceptual Discrimination Tasks to Understand Face Categories and the Cross-Race Recognition Deficit. *Journal of Experimental Psychology: General* 129: 559-574.

Locke, John L. 1993. *The Child's Path to Spoken Language*. Cambridge: Harvard University Press.

Lundqvist, Daniel, Francisco Esteves and Arne Ohman. 1999. The Face of Wrath: Critical Features for Conveying Facial Threat. *Cognition and Emotion* 13: 691-711.

Lyons, John. 1968. *Introduction to Theoretical Linguistics*. Cambridge: Cambridge University Press.

Lyons, John. 1977. *Semantics, 2 vols*. Cambridge: Cambridge University Press.

Markman, Ellen M. 1994. Constraints children place on word meaning. In ed. Paul Bloom, *Language Acquisition: Core Readings*, 154-173. Cambridge: MIT Press.

Matan, Adee and Susan Carey. 2001. Developmental changes within the core of artifact concepts. *Cognition* 78: 1-26.

Matsumoto, Yo. 1993. Japanese numeral classifiers: a study of semantic categories and lexical organization. *Linguistics* 31: 667-713.

Matsumoto, Yo. 1996. Complex predicates in Japanese : a syntactic and semantic study of the notion 'word'. Tokyo: Kurosio Publishers.

McCarthy, John. 1988. Feature Geometry and dependency: A review. *Phonetica* 43: 84-108.
Mehler, Jacques ed. 1983. *Infant and Neonate Cognition*. Hillsdale, NJ: L. Erlbaum Associates.
Mehler, Jacques and Robin Fox eds. 1985. *Neonate cognition : beyond the blooming buzzing confusion*. Hillsdale, NJ: L. Erlbaum Associates.
Merikle, Philip M., Daniel Smilek and John D. Eastwood. 2001. Perception without awareness: perspectives from cognitive psychology. *Cognition* 79: 115-134.
Miller, George A. and Philip N. Johnson-Laird. 1976. *Language and Perception*. Cambridge, MA: Harvard University Press.
Miller, George A. 1998. Nouns in WordNet. In ed. Christiane Fellbaum, WordNet: An Electronic Lexical Database, 23-46. Cambridge: The MIT Press.
Miller, Katherine J. 1998. Modifiers in WordNet. In ed. Christiane Fellbaum, WordNet: An Electronic Lexical Database, 47-67. Cambridge: The MIT Press.
Mohanan, K. P. 1992. Emergence of complexity in phonological development. In eds. Charles A. Ferguson, Lise Menn and Carol Stoel-Gammon, *Phonological Development: Models, Research, Implications*, 635-662. Timonium, Maryland: York Press.
Nazzi, Thierry and Alison Gopnik. 2001. Linguistic and cognitive abilities in infancy: when does language become a tool for categorization? *Cognition* 80: B11-B20.
Needham, Amy. 1999. The role of shape in 4-month-old infants' object segregation. *Infant Behavior & Development* 22: 161–178.
Needham, Amy. 2001. Object recognition and object segregation in 4.5-month-old infants. *Journal of Experimental Child Psychology* 78: 3–24.
Nespor, Marina and Irene Vogel. 1986. *Prosodic Phonology*. Dordrecht, Holland: Foris Publications.
Oliva, Aude and Philippe G. Schyns. 1997. Coarse blobs or fine edges? Evidence that information diagnosticity changes the perception of complex visual stimuli. *Cognitive Psychology* 34: 72–107.
Oliva, Aude and Philippe G. Schyns. 2000. Diagnostic Colors Mediate Scene Recognition. *Cognitive Psychology* 41: 176–210.
Petitto, L. A. and P. F. Marentette. 1991. Babbling in the manual mode: evidence for the ontogeny of language. *Science* 251: 1493-1496.
Petitto, Laura Ann. 1984. *From Gesture to Symbol: the relationship between form and meaning in the acquisition of personal pronouns in American Sign Language*. Unpublished doctoral thesis, Harvard University, Boston, MA.
Petitto, Laura Ann. 1987. On the autonomy of language and gesture: Evidence from the acquisition of personal pronouns in American Sign Language. *Cognition* 27: 1-52.
Petitto, Laura Ann. 1988. 'Language' in the pre-linguistic child. In ed. F. Kessel, *Development of Language and Language Researchers: Essays in honor of Roger Brown*, 187-221. Hillsdale, NJ: Lawrence Erlbaum Associates.
Petitto, L. A., M. Katerelos, B. G. Levy, K. Gauna, K. Tetreault and V. Ferraro. 2001. Bilingual signed and spoken language acquisition from birth: Implications for the mechanisms underlying early bilingual language acquisition. *Journal of Child Language* 28: 453-496.
Pylyshyn, Zenon. 1999.Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences* 22:341–423:
Rehder, Bob and Brian H. Ross. 2001. Abstract Coherent Categories. *Journal of Experimental Psychology-Learning, Memory, and Cognition* 27: 1261–1275.
Rosch, Eleanor H. 1973. Natural categories. *Cognitive Psychology* 4: 328-350.
Rosch, Eleanor. 1975. Cognitive representations of semantic categories. *Journal of Experimental Psychology: General* 104: 192-233.
Rosch, Eleanor, Carolyn B. Mervis, Wayne D. Gray, David M. Johnson and Penny Boyes-Braem. 1976. Basic objects in natural categories. *Cognitive Psychology* 8: 382-439.
Rosenbaum, David A., Richard A. Carlson and Rick O. Gilmore. 2001. Acquisition of intellectual and perceptual-motor skills. *Annual Review: Psychology* 52: 453–70.

Ross, Brian H. and Gregory L. Murphy. 1999. Food for Thought: Cross-Classification and Category Organization in a Complex Real-World Domain. *Cognitive Psychology* 38: 495–553.

Rossell, Susan L., Edward T. Bullmore, Steve C.R. Williams and Anthony S. David. 2001. Brain activation during automatic and controlled processing of semantic relations: a priming experiment using lexical-decision. *Neuropsychologia* 39: 1167–1176.

Sagey, Elizabeth. 1986. *The Representation of Features and Relations in Nonlinear Phonology*. MIT Ph. D. dissertation.

Sagey, Elizabeth. 1988. Degree of closure in complex segments. In eds. Harry van der Hulst and Norval Smith *Features, Segmental Structure and Harmony Processes* Vol. 1. Dordrecht, Foris**:** 169-207.

Schyns, Philippe G. and Luc Rodet. 1997. Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory and Cognition* 23: 681–696.

Schyns, Philippe G. and Aude Oliva. 1999. Dr. Angry and Mr. Smile: When categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition* 69: 243–265.

Schyns, Philippe G. 1997. Categories and percepts: a bi-directional framework for categorization. *Trends in Cognitive Science* 1: 183-189.

Schyns, Philippe G. 1998. Diagnostic recognition: task constraints, object information, and their interactions. *Cognition* 67: 147–179.

Seung, H. Sebastian and Daniel D. Lee. 2000. COGNITION: The Manifold Ways of Perception. *Science* 290: 2268-2269.

Shipley, Thomas F. and Philip J. Kellman. 1997. Spatio-temporal boundary formation: the role of local motion signals in boundary perception. *Vision Research* 37: 1281-1293.

Shipley, Thomas F. 2000. Perception of persistence. In eds. Barbara Landau, John Sabini, John Jonides and Elissa L. Newport, *Perception, Cognition, and Language: Essays in honor of Henry and Lila Gleitman*, 291-309. Cambridge: MIT Press.

Solomon, Karen Olseth and Lawrence W. Barsalou. 2001. Representing Properties Locally. *Cognitive Psychology* 43: 129–169.

Spelke, E. S. and G. A. Van de Walle. 1993. Perceiving and reasoning about objects: insights from infants. In eds. Naomi Eilan, Rosaleen McCarthy and Bill Brewer, *Spatial representation*, Oxford: Basil Blackwell.

Spelke, Elizabeth S. 1982. Perceptual knowledge of objects in infancy. In eds. Jacques Mehler, Edward C.T. Walker and Merrill Garrett, *Perspectives on mental representation : experimental and theoretical studies of cognitive processes and capacities*, Hillsdale, N.J.: L. Erlbaum Associates.

Spelke, Elizabeth S. 1983. Perception of unity, persistence, and identity. In ed. Jacques Mehler, *Infant and Neonate Cognition*.

Spelke, Elizabeth. 1994. Initial knowledge: Six suggestions. *Cognition* 50: 443-447.

Thomas, E., M. M. van Hulle and R. Vogels. 2001. Encoding of categories by noncategory-specific neurons in the inferior temporal cortex. *Journal of Cognitive Neuroscience* 13: 190-200.

Trubetzkoy, N. S. 1939 (1969). *Principles of Phonology*. trans. of *Grundzüge der phonologie* by C. Baltaxe. Berkeley: University of California Press.

Varden, J. Kevin. 2001. *cogPrime: A database of semantic features for investigations into word compositionality*. Tsukuba University Report of the Special Research Project for the Typological Investigation of Languages and Cultures of the East and West Part II. 809-843.

Wei, Longxing. 2001. Lemma congruence checking between languages as an organizing principle in intrasentential codeswitching. International Journal of Bilingualism 5.2:153-173.

Weinreich, Uriel. 1972. *Explorations in Semantic Theory*. The Hague: Mouton.

Werker, Janet and Richard Tees. 1984. Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development* 7: 49-63.

Werker, J. F. 1986. The effect of multilingualism on phonetic perceptual flexibility. *Applied Psycholinguistics* 7: 141-156.

Werker, Janet F. and Judith E. Pegg. 1992. Infant speech perception and phonological acquisition. In eds. Charles A. Ferguson, Lise Menn and Carol Stoel-Gammon, *Phonological Development: Models, Research, Implications*, 285-311. Timonium, Maryland: York Press.

Werker, Janet F. 1994. Cross-language speech perception: Developmental change does not involve loss. In eds. Judith C. Goodman and Howard C. Nusbaum, *The Development of Speech Perception: The transition from speech sounds to spoken words,* 93-120. Cambridge: MIT Press.

Wilson, Martha. 1987. Brain mechanisms in categorical perception. In ed. Stevan Harnad, *Categorical Perception: The groundwork of cognition*, 387-417. Cambridge: Cambridge University Press.

Wyttenbach, Robert A., Michael L. May and Ronald R. Hoy. 1996. Categorical Perception of Sound Frequency by Crickets. *Science* 273: 1542-1544.

Young, Andrew W., Duncan Rowland, Andrew J. Calder, Nancy L. Etcoff, Anil Seth and David I. Perret. 1997. Facial expression megamix: Tests of dimensional and category accounts of emotion recognition. *Cognition* 63: 271-313.

Young, Michael W. 2000. Marking Time for a Kingdom. *Science* 288.5465: 451-453.