What Thoughts Are Made Of

Lera Boroditsky, Stanford University

Jesse Prinz, UNC Chapel Hill

1. Introduction

What are thoughts made of? Do we think in pictures? In words? In symbols?
What is the currency of human cognition and how do the representations that make up
thinking come to be in our minds?   In this chapter we explore the rich sources of input
that humans receive from perception and language and how combining information from
these two input streams can be used to create the amazing complexity and sophistication
of the human knowledge system.

Cognitive science is often seen as emerging from the confluence of two research
programs: Chomsky's nativist critique of behaviorist learning theories and the rise of
artificial intelligence. Together these two tides lead to a seemingly inevitable pair of
conclusions: we think in language-like symbols, and the primitive symbols used in
thought are innate. If we think in innate language-like symbols, then obviously we don't

think in English, or Russian, or Kuuk Thaayorre.   Instead, we think in the universal language of thought - Mentalese (Fodor, 1975).   This conclusion has been explicitly defended by some, in cognitive science, but more often it is an unarticulated background assumption.   In the literature on concepts and categorization, for example, conceptual representations are often described using structured lists of linguistically-labeled features, and researchers rarely suggest that the words used in their theories correspond to mental representations that are radically unlike words.   Nor do they suppose that these words correspond to lexical items in a natural language (as evidenced by the relative lack of cross cultural studies in the first few decades of cognitive scientific research on categories).

In recent times, these assumptions have been critically re-examined and two important (and seemingly contradictory) sources of dissent have emerged.   On the one hand, some critics have argued that we don't think in language-like symbols.   Instead, we think using stored records of sensory and motor states. Following Barsalou (1999), we will call this the perceptual symbols systems hypothesis, or PSS.   On the other hand, some critics have been more willing to accept the claim that language-like symbols are important to thought, but opposed to the claim that the symbols in question belong to an innate and hence universal mental language.   Instead, they venture that statistical regularities found in natural languages (e.g., English, Russian, Kuuk Thaayorre) play an important role in constructing and constituting thought, and that speakers of different natural languages may in fact think in interestingly different ways.   Call this the natural language statistics hypothesis, or NLS.

While these two approaches appear to pull in opposite directions (one away from languaform representations, and one toward them), what they share is a focus on representations being constructed from the inputs, from the patterns an individual observes in the course of their experience in the world (e.g, what they see, what they hear).

In this paper, we discuss the contribution that both of these approaches make to our understanding of how humans construct knowledge, and propose an integration of the two. We will not review the extensive empirical evidence for PSS, since that has been done by other authors in this volume (or see Barsalou et al., 2003). Instead we will discuss what we take to be the major hurdle facing PSS: the problem of abstract ideas. Then we will discuss resources available for coping with this problem and argue that one major resource that has been under-exploited by defenders of PSS is language. We describe ways in which language learning may interact with perceptual symbols and influence cognitive processes. We conclude by drawing some morals about why both perceptual symbols and natural languages are important ingredients in the construction of mature human thought.

2. Perceptual symbols and their limitations

The central idea behind the perceptual symbols systems hypothesis is that the representations (or neural activations) that arise in dedicated input systems during sensation and motor action can be stored and used "offline."   When these stored representations are used in thinking, the brain re-generates a pattern of activation similar to the one that occurred during the perceptual episode. For example, when we access the

idea of a duck, the brain enters an activation state that is like the state we would be in if we were perceiving a duck.

Philosophers in the empiricist tradition have been defending something like PSS for centuries. For example, John Locke (1690) argued against innate knowledge by suggesting that all human concepts can be built simply out of stored copies of experienced perceptual states. Even if humans come equipped with some innate representations, it is clear that a great many representations are learned. This means that we need an account of how concepts can be acquired by means of perception. On the PSS view, perceptual states are simply stored for later use.   On an amodal symbol systems view (such as mentalese), perceptual states can also be stored, but there is an extra step in which an abstract symbol is generated that corresponds to some set of perceptual states. This symbol then (and not the perceptual representations that underlie it) is the representation that is used in thinking. That is, on both views perceptual representations exist in the mind, but on the amodal symbols view, an extra layer of representation (an amodal symbol) is created and used in thinking. The question is, are perceptual representations necessary for conceptual processing, and are they sufficient? And the same for amodal symbols: are they necessary and are they sufficient?

Are perceptual representations necessary in conceptual processing? Consider for example what we know about ducks. A typical list of duck features that make up a duck representation on an amodal symbols view might include things like ducks have feet and feathers and a bill and can swim, etc. This sort of feature list is on first pass appealing as a form of representation because a limited set of features can be used to create many different representations. A feature like "feet" can also be used in other representations

of things that have feet, like humans, pigs, dogs, bathtubs, and so on. Also, the feature list just seems intuitively sensible. Ducks do have feet and a bill and can swim and so on.

But let's consider such feature lists a bit more closely. Imagine that you don't already know what ducks are and are told that ducks have feet and a bill and feathers, and can swim. How would you know that the feet on a duck were not like the feet inside your shoes, or the feet on a pig or a dog or a bathtub? If you don't already know what duck feet look like, simply knowing that ducks have feet would leave open infinite possibilities. Surely ducks don't just have feet, they have <u>duck</u> feet. Further, how would you know where on a duck the feet go? And how many feet? Beyond knowing that ducks have duck feet, you also need to know that they are in number appropriate for a duck and attached like on a duck. And clearly ducks don't just have feathers, they have <u>duck</u> feathers – relatively small and smooth and in a particular duck feather shape, attached like on a duck. And when ducks swim, they don't just swim, they swim like ducks swim. They don't do the backstroke for example. What seemed like a sensible list of features turns out to be vacuous unless one already knows what ducks are like. Every feature in the list must be grounded in perceptual information that an individual already has about ducks, or else it would be useless.

We store a great deal of perceptual information about ducks that is not captured in a typical feature list. We know that duck feet are webbed and that those feet come out of the bottom of a duck, not out of its head. We know the shape of a duck's bill, and we wouldn't mistake a duck bill for a tucan bill or dollar bill. We know that ducks waddle when they walk, that they are larger than doughnuts, and that their eyes are glossy. If we encountered a duck that was even slightly deformed, we would probably notice the

defect. Eyes placed too low, feet too far apart, feathers too sharp—any of these things could be detected. The perceptual information that we store includes shapes (the distinctive curve of a duck's bill) that we would find very difficult to describe.

The perceptual details specific to duck feet and feathers and manners of motion are not simply extra information, they are the essential content of our knowledge that allows us to distinguish a duck from a Las Vegas showgirl, for example, who also likely has feet and feathers and can swim. Without the right perceptual grounding, feature lists are insufficient to account for even simple aspects of human cognition. Of course, the perceptual information we store may be blurry, incomplete, and inaccurate, in various ways, but such as it is, this information is necessary for normal human cognition (such as distinguishing ducks from showgirls, for example).

The argument so far is that perceptual symbol systems are necessary for basic human cognition, and that amodal symbols are insufficient. The next question is whether amodal symbols are necessary to supplement perceptual symbols, or whether perceptual symbols by themselves are sufficient to account for human conceptual ability.

The most persistent and obvious objection to PSS is that it cannot handle abstract concepts. Abstract concepts are, by definition, ones whose category instances are not unified by a shared appearance. According to PSS, human conceptual knowledge is built from stored perceptual states. This may work well for concrete observable physical entities (e.g., ducks and showgirls) that can easily be perceptually experienced. But what about things that we can never see or touch? How do we come to represent and reason about abstract domains like time, justice, or ideas? How do we think about kinship, morality, or politics? Our internal mental lives go far beyond those things observable

through physical experience: we invent sophisticated notions of number and time, we theorize about atoms and invisible forces, and we worry about love, justice, ideas, goals, and principles. How is it possible for the simple building blocks of perception and action to give rise to our ability to reason about domains like mathematics, time, or ideas? The ability to invent and reason about such abstract domains is arguably the very hallmark of human sophistication, so any theory of mental representation worth its salt should have a way of explaining how such abstract notions are acquired and represented.

One strategy for accommodating abstract concepts in a PSS framework is to appeal to metaphor, the idea that abstract domains are understood through analogical extensions from more experience-based domains (e.g., Boroditsky, 2000; Gibbs, 1994; Lakoff & Johnson, 1980).

One of the better-studied examples of such analogical extension is of spatial representations being reused for structuring the more abstract aspects of time. Spatial representations of time abound in our culture, in graphs, time-lines, clocks, sundials, hourglasses, and calendars. In language, time is also heavily related to space, with spatial terms often used to describe the order and duration of events (Clark, 1973; Lakoff & Johnson, 1980; Traugott, 1978). In English, for example, we might move a meeting *forward,* push a deadline *back,* attend a *long* concert or go on a *short* break. Further, people make consistent spatial gestures when talking about time (e.g., Casasanto & Lozano, 2006; Núñez & Sweetser, 2006), with English speakers for example gesturing to the left when speaking about the past and to the right when speaking about the future.

People also appear to spontaneously invoke spatial representations when processing temporal language (e.g., Boroditsky, 2000; Boroditsky & Ramscar, 2002),

such that priming different spatial perspectives will change the way people interpret and process statements about time. People's understanding of time appears so intimately dependent on space, that when people engage in real-life spatial activities such as making an air journey or waiting in a lunch-line, they also unwittingly (and dramatically) change their thinking about time (Boroditsky & Ramscar, 2002). Even simple temporal judgments, it turns out, like being able to reproduce short durations are affected by spatial information (e.g., Casasanto and Boroditsky, 2007). Finally, cultures that rely on different spatial representations, also end up with different representations of time. For example, the Kuuk Thaayorre who think about space in terms of absolute cardinal directions like North, South, East, and West, also lay out time in absolute space – from East to West, unlike English speakers who tend to lay out time from left to right (Boroditsky & Gaby, 2006).

There are of course many other abstract notions to account for beyond time. Some of these have also been linked to spatial representations. For example, people appear to understand kinship concepts spatially; when talking about kin, speakers spontaneously use their hands to draw kinship trees in space (Einfield, 2005). Many other abstract domains have also been shown to elicit consistent spatial gestures, (e.g., up for rising prices, down for falling grades, and so on) (Casasanto & Lozano, 2006). Further, thinking about abstract notions like wealth or honor produces interference for motor actions that are inconsistent with the spatial schemas (e.g., processing a word like "wealthy" makes people slower to make a simple hand movement downward, and processing a word like "poor" makes it harder to make a simple hand movement upward) (Casasanto & Lozano, in press).

A second strategy for accommodating abstract notions in a PSS framework is to appeal to scenarios or scripts (see Shank and Abelson, 1977). For example, the concept of democracy is very hard to visualize because democracies do not look alike. But we certainly know how to act democratically. If asked to settle a problem democratically, we would know to vote, for example. The instructions for doing this can be understood as a script or family of scripts, telling us how to behave. Shank and Abelson thought of scripts as language-like, but they can equally well be implemented by sensory and motor representations of the corresponding behaviors. The democracy script may include representations of hand-raising and other means of casting votes.

A third strategy for accommodating abstract concepts in a PSS framework is to appeal to emotional responses. Consider morality. The range of things we call morally bad have little in common perceptually (stealing, cheating, hitting, and so on). So there cannot be a single image of badness. But all of these things are united by the family of emotions they cause in us. Empirical studies suggest that moral concepts indeed have an emotional basis (Haidt, 2001; Prinz, 2007). For example, hypnotically inducing negative emotions can increase a person's intuition about how wrong something is, even in cases where a described behavior is quite benign (Wheatley and Haidt, 2005). Inducing positive emotions can shift intuitions from a deontological moral frame work (with rules like, thou shalt not kill!) to a consequentialist one (with rules such as, save as many as you can!). Valdesolo and DeSteno (2006) found that people were three times more likely to offer consequentialist responses in moral dilemmas after watching a comedy sketch.

There is also evidence that, when emotions are diminished, people are unable to grasp moral concepts in the normal way. Psychopaths suffer from flattened affect, and

they fail to draw the distinction between moral and conventional rules (Blair, 1995), suggesting that the comprehension of moral rules as such is a matter of emotional responding.   There is also considerable research exploring the specific emotions underlying morality.   Rozin et al. (1999) have showed that crimes against persons elicit anger, crimes against community elicit contempt, and crimes against nature (or "divinity" in non-secular societies) elicit disgust.   Prinz (in prep) has shown that when you yourself perpetrate a crime against a person, the modal response is guilt, but it you perform an unnatural act (e.g., violate a sexual more), the response is shame.   This suggests that there is not one single moral emotion, but many, and these arise in a predictable, context-sensitive way.   This fits in perfectly with PSS.   On that approach, concepts have "variable embodiment" (Barsalou, 1999).   Concepts are temporary constructions in working memory that can vary in their form from context to context.   On the view we are suggesting, the concepts of right and wrong are constituted by a variety of emotions of praise and blame.   On any given occasion, if a person judges that something is wrong, she activates one or another of these emotions.

All three of these strategies (conceptual metaphors, scripts, and emotional responses) show some promise in advancing our understanding of the representation of abstract ideas within a PSS framework.   Of course, much research remains to be done to understand to what extent perceptual information underlies abstract ideas.   While perceptual information may go a long way, there may also be some important limitations. Here, we outline four limitations that strike us as particularly pressing, and then offer a potential solution.

First, some concepts are associated with a very large number of perceptual features, spanning multiple perceptual modalities. In principle, multiple perceptual features can be bound together without any difficulty, but in some cases the binding may become difficult. Consider, for example, some super-ordinate level concepts, such as vehicle. Many vehicles have shared perceptual features (such as wheels or windows), but some are perceptual outliers (a toboggan or hang glider). To keep track of the fact that these belong in the same category as cars, trucks, and boat may be difficult when we are restricted to perceptual features alone.

Second, some concepts may be perceptually grounded in representations that are either structurally complex or temporally protracted, and, when that is the case, there will be a considerable processing cost. Consider the concept of democracy, which we suggested may involve behavioral scripts. It is implausible that the entire script (or family of scripts) runs through the mind every time one thinks about democracy. That would place an exorbitant burden on working memory.

Third, perceptual symbols may also be less than ideal for certain kinds of reasoning processes, especially formal inference. Human reasoning tends to rely on heuristics and biases, and some of these are easy to accommodate within the PSS framework (such as representativeness or availability). But people can also learn to reason in accordance with rules. We can acquire skills for reasoning in accordance with logic for example. We can also reason in domains that make heavy use of symbolic tools, such as mathematics. In math, we can reason about numbers that would be impossible to keep track of perceptually.

A fourth limitation of perceptual symbols is that they are poorly suited for communication. Rich multi-sensory representations are cumbersome to communicate, and even if we use multi-sensory perceptual representations internally, these representations must often be converted to language in order to communicate with others.

These four limitations concern some of the more sophisticated aspects of human cognition: reasoning in terms of broad categories, reasoning about complex systems, or abstract entities, and the ability to communicate complex messages across individuals. While perceptual symbols may get us a lot of the way to complex cognition, there may be domains where perceptual processing alone falls short. Fortunately, humans are not limited to perceptual input produced by the physical world. In addition to the rich perceptual and motor resources we share with other animals, humans also receive a tremendous amount of information through language. In the next section we turn to the role that language plays in shaping and constructing knowledge.

Integrating language into a PSS view of cognition is not hard. Words in a public language are of course also perceived (either through hearing spoken language, seeing signed or written language, or through the tactile channel such as Braille). This means that a stored record of a perceived word or phrase can be treated as a perceptual symbol, or a stored linguistic experience. While the views that thoughts are made of images vs words are often seen as being in opposition, what they share is a focus on the rich sources of information available in the different input channels in human experience. There is a growing body of evidence that people extract and make use of statistical regularities in language. Linguistic perceptual symbols may be very special, in that they may change

the character of thought and extend our cognitive abilities beyond those of creatures that lack language.

3. The Role of Natural Language Statistics

In addition to perceptual and motor experience, people also receive a tremendous amount of information through language. Linguistic input carries highly structured information and comprises many types of statistical regularities. Many of these regularities or patterns of interrelations between sounds, words, and more complex linguistic structures exist only in language, and do not always correspond to the inherent structure of the world. To what extent is people's knowledge constructed out of the patterns of interrelations of elements within the linguistic system?

For example, what role do interrelations between words in a language play in the construction of human knowledge? How much information is extractable simply out of the interrelations of symbols?   And do people actually extract and use such information?

The argument in this part of the paper goes as follows. 1. There is a wealth of information that is extractable out of the internal sets of relations even between entirely ungrounded symbols. Computational models such as LSA and HAL, as well as statistical translation engines (such as the one implemented by Google) capitalize on such information to solve a large array of problems. 2. Humans are also capable of extracting the kinds of statistical patterns of co-occurrence that these computational models rely on, even when the sets of inter-relations are between ungrounded symbols (e.g., novel words with no given referents). 3. When the symbols are grounded (e.g., when the novel words refer to physical objects), people still rely on the patterns of interrelations between the

symbols to inform their representations of the physical objects that the symbols refer to. This is true both for patterns of linguistic interrelations learned in the laboratory, as well as those that exist in natural languages. When patterns of interrelations between words differ across languages, people's representations of those words' referents also differ accordingly. 4. Finally, the patterns of interrelations in language can serve a pivotal role in building representations of abstract entities for which direct perceptual grounding is scant or unavailable. Patterns of linguistic correspondence (such as in conventional metaphor) can guide analogical inference and structure building for abstract domains, such that the perceptually-based knowledge for more concrete domains can be re-used and extended to help reason about abstract entities.

Much information can be extracted just from interrelations of ungrounded symbols. One striking demonstration of how far one can get simply by observing the patterns of interrelations between words (without any knowledge of the physical world) are contextual co-occurrence models such as LSA and HAL (e.g., LSA, Landauer & Dumais, 1997; Landauer, Foltz & Laham, 1998; HAL, Burgess & Lund, 1997). For example, the LSA model is able to extract enough regularity out of the patterns of contextual co-occurrences between words in a large corpus to pass a multiple-choice TOEFL test (Landauer & Dumais, 1997). Such co-occurrence data has also been used to predict an impressive range of semantic priming effects (Lund, Burgess and Atchley, 1995) and provide powerful cues to the syntactic categories of words (e.g., Redington, Chater and Finch, 1998; Burgess and Lund, 1997).

The basic intuition behind such models is that words that occur in similar contexts (that play similar roles in the network of interrelations) will be similar in meaning. If two

words occur in identical contexts then those two words are inter-changeable or synonyms. Extending this logic, the more different the patterns of contextual occurrence between two words, the more different they are likely to be in meaning. Models like LSA work simply by measuring the patterns of contextual co-occurrence between words in a large corpus. LSA gets no perceptual information and has no motor effectors – all of its information comes only from the set of interrelations of ungrounded symbols.

Similar principles can be used to find similarities in structure across complex systems of symbols. For example, the ABSURDIST model (Goldstone, Feng & Rogosky, 2005) is able to find appropriate translations between two non-identical systems by considering only the structure of the similarity relations among the elements within each system, without any external grounding for the elements themselves. Large-scale probabilistic approaches to machine translation based on the same principles (e.g., the Google translation tool: http://www.google.com/translate) have also yielded impressive results.

Of course, just because some information is in principle extractable, as demonstrated by computer simulations, doesn't mean that humans are necessarily capable of or inclined to extract it. Do people in fact extract this kind of information about the interrelations of words in language?

One striking demonstration of how much information is extractable out of language alone comes from the study of color knowledge in people who are congenitally blind (Shepard & Cooper, 1992). Shepard and Cooper asked people to rate the similarity of colors to one another, for example, how similar violet is to red, or orange to gold, or green to yellow, etc. When people with normal color vision are asked to do this, their set

of color similarity ratings is best captured by a color circle where adjacent colors on the wheel are those that are most similar to one another (e.g., yellow is next to gold, violet is next to purple, etc), and colors that are seen as very dissimilar are on opposite sides of the circle (e.g., green is opposite red). But what if the same questions were posed to people who have never perceptually experienced color, people who are congenitally blind? While the shape of the color space that emerges from the similarity ratings of the congenitally blind is very different in overall shape, what is striking is how many short-distance relationships are preserved. For example, purple is rated as most similar to violet, yellow as most similar to gold, orange to red, and so on. This information about colors is obviously only available through language to people who are congenitally blind. It is remarkable how much information is extracted out of the patterns in language in the absence of any supporting perceptual information.

Many other studies have revealed the exquisite sensitivity that humans have for extracting complex statistical structure out of linguistic input (e.g., Saffran, Aslin & Newport, 1996; Gomez & Gerkin, 1999; Saffran, 2002). For example, adults and even very young infants can use sound co-occurrence patterns in speech to do things like discover word boundaries and abstract grammatical patterns. Importantly for the purpose of this paper, people's understanding of linguistic content (i.e., word and utterance meaning) is also affected by patterns of word co-occurrence (Boroditsky & Ramscar, 2003).

For example, Boroditsky & Ramscar (2003) exposed people to a stream composed entirely of novel words. In this stream of novel words, all critical words occurred equally often with each other, but the patterns of contextual occurrence differed.

Some words occurred in the same contexts, and some did not. For example, the words FAP and ZUN never co-occurred with one another, but each occurred frequently with the word REC. Another set of words like KIF and DUT would never co-occur with REC, but would instead co-occur with NIZ, and so on. The subjects were told they were spying on an alien mobile-phone conversation and were asked to learn as much about the meanings of the alien words as they could. When asked to rate the meaning similarity of pairs of novel words, subjects rated those words that occurred in the same contexts (e.g., FAP and ZUN in the example above) as being more similar in meaning.

It appears that humans (much like the simple mechanisms underlying co-occurrence models of semantics such as LSA) do pick up on contextual co-occurrences of novel words, and that they allow this statistical information to inform their judgments about the words' meanings. Of course, participants in this study had no other information about the words' meanings except the patterns of contextual co-occurrence. A further question is whether people would allow co-occurrence statistics in language to inform their similarity judgments about more contentful representations. For example, can co-occurrence relationships between words affect people's representations of those words' referents?

To investigate this question, Boroditsky & Ramscar (2003) taught people nonce names for a set of novel objects. After learning the objects' names, participants then witnessed an alien "conversation" about those objects. In the conversation, the names of the objects co-occurred with context words in the same way as described above. Later subjects rated the similarity of the objects themselves (using their pictures). The results showed that the co-occurrence properties of the objects' names influenced people's

perceptions of similarity of the objects themselves. Pictures of objects whose nonce names occurred in the same nonce linguistic contexts were judged to be more similar to one another than pictures of objects whose names did not share linguistic contexts. It appears that people not only extract patterns of co-occurrence from the linguistic stream and use the patterns of co-occurrence to inform judgments about the meanings of words, they also allow this information from the linguistic stream to influence their judgments about the words' referents (in this case judgments about the similarity of pictures of objects).

These results suggest that there may be an interesting bi-directional relationship between linguistic co-occurrence and representation. Objects whose names occur in the same linguistic contexts come to be perceived as being more similar. Of course as objects become more similar, it is more likely that they will be talked about in ever more similar contextual collocations, which might serve in turn to further strengthen the similarities between the objects themselves and so on. In this way, it appears that co-occurrence statistics may offer people a cognitive mechanism that allows them to sharpen and refine their underlying conceptual representations over time.

A further question to ask is whether people's ability and willingness to learn and use patterns of interrelations between words extends beyond small-scale laboratory demonstrations. The co-occurrence patterns in natural language are far more complex than those introduced in the lab, and need to be computed over thousands of word types and many millions of word tokens. Do people naturally extract patterns of interrelations from the giant sea of information they are exposed to in the course of normal language use?

The fact that patterns of interrelations between words differ across languages provides a nice natural test-bed for this question. If people allow patterns of interrelations between words to influence their mental representations of the words' referents, then when patterns of interrelations between words differ across languages, people's representations of those words' referents should also differ.

One natural language example of this idea can be found in the domain of grammatical gender. In languages like Spanish and French, all nouns are divided into either the masculine or feminine grammatical gender. Nouns that are grammatically masculine will co-occur with the same masculine articles, pronouns, adjective endings and so on that are used to talk about biological males. Nouns that are grammatically feminine will co-occur with the grammatical markers used for talking about biological females. Conveniently for our purposes, grammatical gender assignments of nouns differ across languages. For example, the word for the sun is grammatically feminine in German but grammatically masculine in Spanish. This allows us to ask whether the relations words bear to each other within a linguistic system (in this case by belonging to the same grammatical category and by virtue of that sharing patterns of contextual co-occurrence) has an influence on how users of that linguistic system conceive of the words' referents. Do German speakers think of the sun as being more like a biological female than do Spanish speakers?

It turns out that such patterns of interrelations between words do influence people's representations of objects in the world. When an object's or entity's name is grammatically masculine in a language, speakers of that language will describe that object using more masculine adjectives, will rate the object as being more similar to

biological males, will rate the objects as having more masculine properties, will be more likely to personify the object or entity with a masculine voice and/or body, and so on (e.g., Boroditsky et al, 2003; Sera et al, 1994; Jacobson, 1959).

Similar findings have been obtained with other types of relations between words. For example, objects whose names share morphological roots in Dutch but not English will be rated to be more similar by Dutch speakers. For example, in Dutch the words for sink (wasbak) and garbage can (vuilnisbak) share the root of –bak, and as a result have phonological similarity. In English, the names for these objects do not have this kind of relationship. Correspondingly, Dutch speakers judge pictures of sinks and garbage cans to be more similar to one another than do English speakers (Baayen & Boroditsky, 2004).

Cross-linguistic differences have been found even in very basic perceptual domains like color. For example, English and Russian color terms divide the color spectrum differently. Unlike English, Russian makes an obligatory distinction between lighter blues ('goluboy') and darker blues ('siniy'), and this linguistic difference leads to differences in color discrimination (Winawer et al, 2007). Russian speakers are faster to discriminate two colors if they fall into different linguistic categories in Russian (one siniy and the other goluboy) than if they are from the same linguistic category (both siniy or both goluboy). English speakers tested on the same stimuli show no such differences across the goluboy-siniiy border.

In sum, there's a growing body of evidence that people extract complex statistical regularities out of language, and further that they incorporate these extracted patterns into their knowledge representations.

Further, there is evidence that patterns of interrelations within language can help us construct new representations by encouraging analogical or metaphorical extensions between knowledge domains that share linguistic descriptions (see discussion of abstract concepts above). For example, talking about time using spatial terms encourages people to reuse spatial representations for the purpose of thinking about time. Using spatial representations for thinking about time may allow us to build representations of time that go beyond our physical experience. For example, we often talk about time as if it were a physical path (e.g., we're approaching the holidays, we're getting close to the deadline, we're coming up on New Years, etc). Once the metaphor is in place and time is conceived of as a path, new possibilities open up for thinking about time that go beyond our normal experience. For example, a physical path you can travel in any direction you want, and at whatever speed you want. Extending this idea to time, gives us the notion of time-travel, not something we have any personal physical experience with. In this way, abstract (or maybe even physically impossible) notions like time-travel can be built out of concrete representations by creating and extending metaphors between domains of knowledge.

Further, there is evidence that language plays a constructive role in this process, directing which representations in our spatial knowledge will be reused for thinking about time. While it may be a universal that spatial representations are used for time, languages and cultures differ in terms of how time is laid out in space. For example, Núñez & Sweetser (2006) observed that the Aymara talk about the future as being behind them and the past as being ahead of them, and gesture accordingly. English and Mandarin differ in terms of how often they talk and think about time vertically, with Mandarin speakers

being much more likely to use vertical metaphors for time than do English speakers (e.g., Boroditsky, 2001).

Of course, language can play this constructive role in knowledge building not just in building abstract domains, but also in expanding our knowledge of the physical world well beyond what we experience in our own lifetimes. Once a number of experienced perceptual states are stored, those representations can be used in new combinations to conjure up distant lands we have yet to visit, or times long ago we have no chance to experience. This conjuring is often done for us through language, by using words in new combinations to create new combinations of perceptual states.

While ungrounded linguistic symbols alone have clear limitations (see our discussion of ducks and showgirls earlier), when combined with some stored perceptual information linguistic symbols can be used to greatly expand our knowledge. Even in mundane domains, combining a bit of perceptual information with the combinatorial power of language can be extremely useful. For example, imagine you have perceptual information stored about grapefruits and other citrus fruits, but you have never seen or heard of a pomelo. If you are told that pomelos are citrus fruits similar to grapefruits, but larger and sweeter and with a thicker rind, you can easily re-use the perceptual information you have stored about other citrus fruits to form a pretty good representation of pomelos without ever having had experience with one. Given this linguistic description and the perceptual information you have stored about other citrus fruits, you would probably be able to identify a pomelo at a fruit-stand, and can make some good predictions about how pomelos might behave, look, taste, sound, smell, feel in the hand, and so on.

Boroditsky   7/25/07   23

It is possible that much of our mental representation of the physical world is in fact constituted not out of direct experience, but out of re-used perceptual representations, with the re-use guided by what we hear in language. This allows us to build knowledge not just out of what we ourselves experience in the physical world, or what we ourselves can conjure up out of our own experiences, but also out of representations others lead us to create, conjured up through language.

4. Summary

Neither perceptual information alone, nor the sets of correspondences between elements in language alone are likely to be able to amount to the sophistication, scale, and flexibility of the human conceptual system. Luckily, humans receive heaping helpings of both of these types of information. Combining information from these two input streams, as well as extracting the wealth of information that exists in the correspondences across input streams can help overcome the shortcomings of relying on any single information stream and can reveal information not available in any one stream. The possibilities for building representations improve drastically when both external grounding information and sets of system-internal interrelations are allowed to play a role.

The findings reviewed in this paper begin to tell a story about the acquisition of knowledge that is simultaneously grounded in perceptual experience and enmeshed within a linguistic system. Conceptual knowledge does not appear to be removed from its perceptual origins, nor from its linguistic origins.   It seems that human knowledge is built largely from experience, from the inputs, from observed statistical regularities in the

physical world and in language. Appreciating the richness of the information humans

receive as input both from these two sources should help us turn what had seemed to be

unsolvable mysteries about the origins of knowledge into tractable puzzles.

References:

Baayen, R.H., & Boroditsky, L. (2004, October 8). Effects of lexical structure on picture

    comparison. Edmonton, Conference: Conceptual Structure, Discourse &

    language.

Barsalou, L.W. (1999). Perceptual symbol systems. Behavioral and Brain Sciences, 22,

    577-609.

Barsalou, L.W., Simmons, W.K., Barbey, A., & Wilson, C.D. (2003). Grounding

    conceptual knowledge in modality-specific systems. Trends in Cognitive

    Sciences, 7, 84-91.

Blair, R. (1995). A cognitive developmental approach to morality: investigating the

    psychopath. Cognition, 57, 1-29.

Boroditsky, L. (2000). Metaphoric Structuring: Understanding time through

    spatial metaphors. Cognition, 75(1), 1-28.

Boroditsky, L. (2001). Does language shape thought? English and Mandarin

    speakers' conceptions of time. Cognitive Psychology, 43(1), 1-22.

Boroditsky, L. & Ramscar, M. (2002). The Roles of Body and Mind in Abstract

    Thought. Psychological Science, 13(2), 185-188.

Boroditsky, L., Schmidt, L. A., & Phillips, W. (2003). Sex, syntax and semantics.

    In D. Gentner & S. Goldin-Meadow (Eds.), Language in mind: Advances

    in the study of language and thought (pp. 61-78). Cambridge, MA: The

    MIT Press.

Boroditsky, L. & Ramscar, M. (2003). Guilt by association: Gleaning meaning

from contextual co-occurrence. Proceedings of the 25th Annual Meeting

of the Cognitive Science Society, Boston, MA.

Boroditsky, L, & Gaby, A. (2006). East of Tuesday: Representing time in

absolute space. In the Proceedings of the 28th Annual Meeting of the

Cognitive Science Society, Vancouver, Canada.

Burgess, C., and Lund, K. (1997). Modeling Parsing Constraints with High-Dimensional

Context Space. Language and Cognitive Processes, 12, 177-210.

Casasanto, D. & Boroditsky, L. (2007). Time in the Mind: Using space to think

about time. Cognition. doi:10.1016/j.cognition.2007.03.004

Casasanto, D. & Lozano, S. (2006). Metaphor in the Mind and Hands. Proceedings of

28th Annual Conference of the Cognitive Science Society. Vancouver, BC. 142-

147.

Casasanto, D. & Lozano, S. (in press). Meaning and Motor Action. Proceedings of 29th

Annual Conference of the Cognitive Science Society. Nashville, TN. (To appear

August, 2007)

Clark, H. H. (1973). Space, time, semantics, and the child. In T. E. Moore, Cognitive

development and the acquisition of language, New York: Academic Press.

Enfield, N. J. 2005. The body as a cognitive artifact in kinship representations. Hand

gesture diagrams by speakers of Lao. Current Anthropology, 46, 51-81.

Fodor, J. (1975). The Language of Thought. Cambridge, MA: Harvard University Press.

Gibbs, R.J. (1994). The poetics of mind: Figurative thought, language, and understanding. New York, Cambridge University Press.

Goldstone, R. L., Feng, Y., & Rogosky, B.. Connecting to the world and each other. In D. Pecher & R. Zwaan (Eds.) *Grounding cognition: The role of perception and action in memory, language, and thinking*. Cambridge: Cambridge University Press. (pp. 292-314)

Gomez, R.L., & Gerken, L.A. (1999). Artificial grammar learning by one-year-olds leads to specific and abstract knowledge. Cognition, 70, 109-135

Haidt, J. (2001). The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment. Psychological Review, Vol. 108. No. 4, 814-834.

Jakobson, R. (1959). On linguistic aspects of translation. In R.A. Brower (Ed.), On translation (pp. 232-239). Cambridge, MA: Harvard University Press.

Lakoff, G., & Johnson, M. (1980). Metaphors we live by. Chicago: University of Chicago Press.

Landauer, T.K., and Dumais, S.T. (1997). A Solution to Plato's Problem: The Latent Semantic Analysis Theory of Acquisition, Induction, and Representation of Knowledge. Psychological Review, 104, 211-40.

Landauer, T.K., Foltz, P.W., and Laham, D. (1998). An Introduction to Latent Semantic Analysis. Discourse Processes, 25, 259-84.

Locke, J. (1690/1979). An essay concerning human understanding. P. H. Nidditch, ed. Oxford: Oxford University Press.

Lund, K., Burgess, C., and Atchley, R.A. (1995). Semantic and Associative Priming in High-Dimensional Semantic Space. Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society. Pittsburgh, PA: Erlbaum. 660-65.

Núñez, R.E., & Sweeser, E. (2006). With the Future Behind Them: Convergent Evidence From Aymara Language and Gesture in the Crosslinguistic Comparison of Spatial Construals of Time, Cognitive Science, 30(3), 401-450.

Prinz, J. J. (2002). Furnishing the mind: Concepts and their perceptual basis. Cambridge, MA: MIT Press.

Prinz, J. J. (2007). The Emotional Construction of Morals. Oxford: Oxford University Press.

Redington, M., Chater, N., and Finch, S. (1998). Distributional Information: A Powerful Cue for Acquiring Syntactic Categories, Cognitive Science, 22(4), 425-469.

Rozin, P., Lowery, L., Imada, S. and Haidt, J. (1999). The CAD triad hypothesis: A mapping between three moral emotions (contempt, anger, disgust) and three moral codes (community, autonomy, divinity). Journal of Personality and Social Psychology 76: 574–86.

Saffran, J.R., Aslin, R.N., & Newport, E.L. (1996). Statistical learning by 8-month old infants. Science, 274, 1926-1928.

Saffran, J. R. (2002). Constraints on statistical language learning. Journal of Memory and Language, 47, 172-196.

Sera, M. D., Berge, C., & del Castillo, J. (1994). Grammatical and conceptual forces in the attribution of gender by English and Spanish speakers. Cognitive Development, 9(3), 261-292.

Shepard, R. N. Cooper, L. A. . Representation of colors in the blind, color--blind, and
    normally sighted. Psychological Science, 3(2), 97--104.

Shank, R. C., & Abelson, R. (1977). Scripts, plans, goals and understanding. Hillsdale,
    NJ: Elbarum.

Traugott, E. C. (1978). In J. H. Greenberg, On the expression of spatiotemporal relations
    in language, Word structure, Universals of human language, vol. 3. (pp.
    369±400). Stanford, CA: Stanford University Press.

Valdesolo, P., & DeSteno, D. (2006).Manipulations of Emotional Context Shape Moral
    Judgment. Psychological Science, 17 (6), 476–477.

Wheatley, T. & Haidt, J. (2005).Hypnotic Disgust Makes Moral Judgments More Severe.
    Psychological Science, 16 (10), 780–784.

Winawer, J., Witthoft, N., Frank, M., Wu, L., Wade, A., and Boroditsky, L.
    (2007). Russian blues reveal effects of language on color discrimination.
    PNAS published April 30, 2007, 10.1073/pnas.0701644104