



How Language Programs the Mind

Gary Lupyan,^a Benjamin Bergen^b

^a*Department of Psychology, University of Wisconsin-Madison*

^b*Department of Cognitive Science, University of California at San Diego*

Received 21 February 2014; received in revised form 31 October 2014; accepted 31 October 2014

Abstract

Many animals can be trained to perform novel tasks. People, too, can be trained, but sometime in early childhood people transition from being trainable to something qualitatively more powerful—being *programmable*. We argue that such programmability constitutes a leap in the way that organisms learn, interact, and transmit knowledge, and that what facilitates or enables this programmability is the learning and use of language. We then examine *how* language programs the mind and argue that it does so through the manipulation of embodied, sensorimotor representations. The role language plays in controlling mental representations offers important insights for understanding its origin and evolution.

Keywords: Language; Language and thought; Embodiment; Categorization; Language evolution

1. People are programmable

In Surat Thani, a province in Southern Thailand, there is a school where pigtail macaques are trained to gather coconuts. The monkeys clamber up palms rising up to 80 feet, and detach the ripe coconuts by spinning them off their stalks. Those who complete 3–6 months of elementary training can proceed to “secondary school” and “college,” where they are trained by experienced human trainers to use more advanced gathering techniques. A successful graduate can gather up to 1,500 coconuts a day. But the process by which the monkeys reach this impressive state is arduous. First, they have to be coaxed to simply touch a coconut, then gradually to start spinning it. Each behavior is shaped by rewards, bit by torturous bit, until the desired task-set is reached months later. Similarly laborious processes are necessary in training animals to perform far simpler tasks,

Correspondence should be sent to Gary Lupyan, Department of Psychology, University of Wisconsin-Madison, 1202 W Johnson Street, Madison, WI 53706. E-mail: lupyan@wisc.edu

whether teaching a pet dog to sit on command or a laboratory animal to press a button when it sees a circle (Harlow, 1959).

Like other animals, humans can be trained. But around the age of 2, humans transition from being merely trainable to something qualitatively more powerful—being *programmable*. We can sculpt the minds of others into arbitrary configurations through a set of instructions, without having to go through laborious trial-and-error learning. We can cause someone to imagine something, to recall a memory, to do (or not do) something. We can also turn these instructions inwards and use them to control our own minds (Rumelhart, Smolensky, McClelland, & Hinton, 1986; Vygotsky, 1962). Many animals can be *trained* to press a button any time they see a circle. Humans can simply be *told* to do it, the instruction bypassing the need for slow, incremental adjustments made by trial-and-error learning, and instead—in an instant—reprogramming the listener’s mind into a task-relevant state.

We argue that the capacity for programming constitutes a leap in the way that organisms learn, interact, and transmit knowledge, and that language is a key player in the emergence of this capacity. The capacity for language that has emerged in the human lineage can be understood as a powerful system for flexibly shaping learning and behavior in adaptive ways (Hays, 2000) both across and within generations of learners.

Language affects programmability on two timescales. At the longer timescale, learning language qualitatively alters the learners’ conceptual space and helps to align the conceptual repertoires of speakers within a speech community. At the shorter timescale, language augments categorization, highlighting task-relevant dimensions (in effect, collapsing analogue high-dimensional mental spaces to lower-dimensional, more discretized ones) (Lupyan, 2012).

We first describe some functions that programmability serves in human cognition, and point to the crucial involvement of language in each. We next examine *how* language programs the mind and argue that it does so through the manipulation of embodied, sensorimotor representations. Language acts directly on mental states (Elman, 2009). We then examine what design features of words—a key component of language—make them especially well suited to programming the mind. Finally, we relate this framing to the question of language evolution, and suggest that programming minds may have served as a major pressure in the emergence of language.

2. Four examples of programmability, and the role of language in each

2.1. Teaching

A student monkey’s coconut-gathering behavior is shaped bit by bit because there is no alternative; no shortcut. In contrast, humans can be told about the goal in seconds. They might not be particularly good at accomplishing the goal at first and they will never be as adept as the monkeys at climbing trees, but they will know from the very start what is being demanded of them. Communicating the goal state—perhaps the most laborious

part of training the macaques—can be accomplished almost instantaneously in the case of humans. One does not have to look far to find other examples (adaptive even in the strictest Darwinian sense). Consider the simple routine of looking both ways when crossing a busy street—a domain ill suited to trial-and-error learning. Developing this routine as an automatic habitual behavior takes practice—a kind of training which *can* take a long time. But in humans, getting the behavior off the ground and establishing the objective can be programmed with a few simple words (“Look both ways before crossing the street”).

Either of these behaviors—collecting coconuts or crossing streets—could in principle be learned through imitation or non-verbal tutoring.¹ In actual practice though, they are often learned through instruction in the form of verbal directives. In part, this is because there are things linguistic instruction can do that imitation cannot. Even in learning nuanced sensorimotor tasks like playing the violin or learning to fly a plane,² we can (and do) use language to guide the learning process, to “tame the path-dependence” (Clark, 1998), choosing to present one set of examples or experiences to the learner over another, and to nudge the student’s actions, often with verbal commands (“do that once more,” “no, not quite so hard,” “look at what I am doing right now”). Verbal instruction is a particularly effective method of eliminating sources of error (Harlow, 1959)—a primary cause of the difficulty of training non-human animals.

Verbal instructions can also be turned inward, guiding one’s own behavior—an idea most closely associated with Vygotsky (1962)—but one that has since received support using more rigorous testing methodologies. For example, labeling one’s actions supports the integration of event representations (Karbach, Kray, & Hommel, 2011), overt self-directed speech can improve performance on such tasks as visual search by helping to activate visual properties of the targets (Lupyan & Swingley, 2012). Conversely, interfering with (covert) verbalization impairs the ability to flexibly switch from one task to another (Baddeley, Chincotta, & Adlam, 2001; Emerson & Miyake, 2003), hinting that normal task-switching performance is aided by such self-directed, covert instruction.

2.2. Category learning and category alignment

Communicating, teaching, and coordinating action is easier when members of a community operate in a shared conceptual space. So, in instructing a child to look both ways, it helps to have shared concepts of the directional terms *left* and *right*. Where do these shared concepts come from? Certainly, one source is the common evolutionary history of humans that gave rise to similar bodies and similar sensory-motor capacities. However, we argue, following Whorf (1956; see also Cassirer, 1962), that our shared biology and environment drastically under-determine the possible concepts humans may end up with (see Malt et al., 2015 for recent discussion). A common language helps to create a shared conceptual space, which is critical for programmability (and interoperability). How? A minimum in learning many of the same words speakers of a given language are all guided to become practiced in making the same categorical distinctions (Levinson, 1997). This idea that language helps to create and align concepts conflicts with the still-popular idea that

words simply map onto a pre-existing conceptual space ([Bloom, 2001](#); [Li & Gleitman, 2002](#); [Snedeker & Gleitman, 2004](#)) according to which language is simply a tool for communicating our universally shared ideas. But mounting evidence has made this position a difficult one to sustain (see also [Gomila, 2011](#); [Lupyan, 2012](#) for discussion).

It is easiest to see how language can drive conceptual development in cases where individuals have access to language but lack access to its referents. Consider color concepts in people who are congenitally blind. Although they lack sensory experience with color altogether, such individuals still develop a color vocabulary and reveal a semantic color similarity space very close to that of sighted people ([Marmor, 1978](#)). For instance, a blind person can report that the sky is typically blue and that purple is more similar to blue than to green. This is only possible because blind people, like sighted people, are exposed to statistical regularities through the use of color words produced by their language community. The concept of blue that a blind person has is certainly different from the concept of a sighted person, but without language the blind person would lack all knowledge of color altogether!

This argument is not simply rhetorical. Growing evidence suggests that labels (category names) play a causal role in category learning in children (e.g., [Casasola, 2005](#); [Nazzi & Gopnik, 2001](#); [Plunkett, Hu, & Cohen, 2008](#); [Waxman & Markow, 1995](#); [Althaus & Mareschal, 2014](#); but see [Robinson, Best, Deng, & Sloutsky, 2012](#)) as well as adults ([Lupyan & Casasanto, 2014](#); [Lupyan & Mirman, 2013](#); [Lupyan, Rakison, & McClelland, 2007](#)). The case for the influence of language on concept learning is stronger still for concepts whose members do not cohere on any perceptual features, such as the concept of exactly EIGHTEEN ([Frank, Everett, Fedorenko, & Gibson, 2008](#); [Gordon, 2004](#))³ Thus, words not only carve nature at its joints, but they help to carve joints in nature, reifying conceptual distinctions.

The impact of language on the construction of aligned conceptual categories happens on two timescales. First, in the course of learning the words of their language, individuals are led to learn a set of concepts *and* align those concepts to a much greater degree than they would otherwise. Subsequently, language use allows people to program one another (as well as to program themselves). Learning a word—which is often a protracted, procedural process (but see [McMurray, Horst, & Samuelson, 2012](#) for a fuller account)—involves learning to instantaneously activate those representations using words as cues. To take one example, learning color names involves partitioning the color space into nameable regions—often a slow process ([Wagner, Dobkins, & Barner, 2013](#)). But once learned, a color term can rapidly instantiate a categorical color representation allowing one person to tell another to look for red berries, not the blue ones—instantly focusing one’s attention on the dimension of interest.⁴ The next two sections develop this idea further.

2.3. *Learning and using abstract concepts*

William James wrote that the power to evaluate sameness “is the very keel and backbone of our thinking” ([James, 1890](#), p. 459). Sameness matters because while no two entities or experiences are ever exactly the same, we must nevertheless treat them as the

same lest we suffer the same fate as Borges's *Funes* who, perceiving each event as a distinct entity, was "disturbed that a dog at three-fourteen (seen in profile) should have the same name as the dog at three-fifteen (seen from the front)" (Borges, 1942/1999, p. 136). Without the ability to evaluate sameness, there are no categories (indeed, one can think of categorization as the representation of sameness in the given context (Harnad, 2005; Lupyan, Mirman, Hamilton, & Thompson-Schill, 2012)).

We do not need language for learning to categorize two views of a dog as both being of a dog (though see Collins & Curby, 2013), but language nevertheless plays an important and sometimes critical role in learning to selectively represent items in a way that promotes their categorization (Lupyan, 2009; Lupyan & Mirman, 2013; Lupyan & Spivey, 2010). Consider the category of the sameness relation itself. You are shown two pictures. If they are the same, press button A. If they are different in some way, press button B. A categorization task like this involves abstracting over everything except the *relation* between the two items. This kind of task is trivially simple for adult humans who are appropriately instructed (Farell, 1985). One might be tempted to conclude that people have a special knack for discovering abstract relations (Penn, Holyoak, & Povinelli, 2008) or come innately equipped with concepts like "same" and "different" (Mandler, 2004). But it turns out that this task becomes considerably more difficult when the benefit of being programmed with the words "same" and "different" is withheld. When otherwise competent adults are left to the vicissitudes of trial and error learning, the task suddenly becomes challenging to the point where many adults fail to infer these suddenly not-so-basic relations (Castro & Wasserman, 2013; Young & Wasserman, 2001). Young children are also frequently flummoxed (Addyman & Mareschal, 2010; but see Walker & Gopnik, 2014 for successful reasoning about relational categories by young children in a very different task). Although both children and non-human animals can rely on perceptual cues to use the amount of visual disorder (entropy) as a measure of difference, apparently only older children and adults are able to transcend this dependence on perceptual cues, successfully generalizing the relation to two-item same/different displays.⁵

On the account proposed here, language plays a role in learning the same/different relation in the following way. First, learning words like "same" and "different" *necessitates* learning the appropriate concepts. Although humans and perhaps some other animals are *capable* of learning to categorize using the same/different relation, the active use of the words *same* and *different* provides language-learners with the structured experience that ensures the learning takes place. We do not claim that it is impossible to learn such relations without the accompanying words. Rather, the guidance provided by word-learning ensures that certain learning episodes take place. Episodes that might not occur in the absence of the linguistic guidance (Ozcaliskan, Goldin-Meadow, Gentner, & Mylander, 2009 for related discussion). This first phase of programmability lays down the procedures that facilitate or enable participants to posit and test hypotheses such as "Is this about the two pictures being the same or not?" But judging by the variability which even linguistically competent adults show on the task in the absence of explicit instruction, many people either do not posit, or do not successfully test this hypothesis. The second part of our language-based programmability perspective is that words are used to program the mind

on-line: rapidly and reliably configuring the appropriate task-set (i.e., collapsing the hypothesis space) and categorizing the inputs into the appropriate relational categories.

2.4. *Invoking the familiar, the new, and the hypothetical*

We now know that imagination shares much with perception; *imagining* a red circle is similar in many respects to the mechanisms that underlie *seeing* a red circle (Kosslyn, Ganis, & Thompson, 2001). Activation of visual representations in perception requires the appropriate visual input. Activation of the (largely overlapping) representations in imagery requires a cue that can invoke this visual representation in the absence of the visual input. Where do these cues come from? We can construct an environment in which certain perceptual experiences (e.g., shapes) become associated with certain other perceptual experiences (e.g., certain sounds). The sound can then cue perceptual representations of the associated shape. An animal trained in this way can be cued to “visualize” a certain shape by being presented with the appropriate sound. Scale this up and one begins to see the basic outlines of language. We are not advocating a return to the Skinnerian idea of language-learning as a form of conditioning. We are simply pointing out that in its stripped-down, basic form, learning a language provides users with a very large set of mutually comprehensible cues that can be used to control, to “sculpt” mental representations including low-level perceptual ones (Deak, 2003; Hays, 2000; [Lupyan & Ward, 2013](#)).

As remarkable as it is that we can invoke a perceptual representation in someone simply by uttering a few sounds, it is even more remarkable is that we can cue (read: program) people to construct entirely novel representations simply by verbalizing an instruction. We can, for example, ask someone to “Draw a person”—an instruction that presumably activates some visual representation of a goal state and a sequence of motor movements to help achieve that goal state. But we can also ask someone to “Draw a person that does not exist” ([Karmiloff-Smith, 1990](#)). Children younger than 5 years struggle with the latter task, tending to exaggerate elements within the familiar schema. Older children are able to incorporate additional elements (an extra head, wings). Further empirical work is needed to know whether this ability to improvise on familiar schemas, to go beyond the real, is aided by linguistic control.

3. Language programs the mind by interfacing with grounded mental states

In the previous section, we outlined some ways in which language is implicated in programming the mind. We now consider two broad possibilities for how this is accomplished. The first is that the mind has an inventory of amodal and abstract concepts derived from evolution and experience (Mentalese; a “language of thought”). On this view, natural language can push these representations around, but the representations that constitute knowledge reside in a separate store to which language does not productively contribute ([Bloom & Keil, 2001](#); [Fodor, 2010](#)). On the second view, learning and using

language directly modulates the very same mental states that constitute perception and action (e.g., [Elman, 2009](#); [Lupyan, 2012](#)). On this view, there is no need for a language of thought. It is not that we think “in” language. Rather, language directly interfaces with the mental representations that are used in perception and action, helping to form the (approximately) compositional, abstract representations that thinkers like Fodor take as a priori. We believe the empirical evidence supports this second view.

In the last several decades there has accumulated a substantial body of research, often subsumed under the label “embodied” or “grounded” cognition. One of the central claims of the linguistic branch of this research program is that word meanings are grounded in sensorimotor experiences. Producing and comprehending language involves automatic engagement of neural systems that reflect the perceptual, motor, and affective features of the content of the language (see [Bergen, 2012](#), for review). These effects are variably called perceptuomotor traces ([Zwaan, Madden, Yaxley, & Aveyard, 2004](#)), simulations ([Barsalou, 1999](#)), or resonances ([Taylor & Zwaan, 2008](#)). But the recurring empirical result is that understanding a word involves, to varying extent, representing the real-world referents of that word ([Evans, 2009](#); [Kiefer & Pulvermüller, 2012](#)). For example, comprehending a word like “eagle” activates visual circuits that capture the implied shape ([Zwaan, Stanfield, & Yaxley, 2002](#)), canonical location ([Estes, Verges, & Barsalou, 2008](#)), and other visual properties of the object, as well as auditory information about its canonical sound ([Winter & Bergen, 2012](#)). Words denoting actions like *stumble* engage motor, haptic, and affective circuits ([Glenberg & Kaschak, 2002](#)).

A reasonable question often raised about these findings is whether the motor, perceptual, and affective representations activated by words is epiphenomenal—perhaps words tend to activate these representations downstream but word *meanings* involve a totally distinct set of mechanisms ([Mahon & Caramazza, 2008](#)). However, there is now extensive evidence that not only are perceptual, motor, and affective systems activated during meaning construction, but that this activity plays a functional role in comprehension. We know this in part from findings that perceptual and motor information comes online very early—as quickly as 100 ms from word onset ([Amsel, Urbach, & Kutas, 2014](#); [Pulvermüller, Shtyrov, & Hauk, 2009](#)). The causal function of these activations is further supported by studies showing that disrupting perceptual or motor resources decreases comprehension speed ([Glenberg et al., 2010](#); [Pulvermüller, Shtyrov, & Ilmoniemi, 2005](#)) and accuracy ([Yee, Chrysikou, Hoffman, & Thompson-Schill, 2013](#)) of words whose meanings involve those specific percepts or actions.

Critically, as hinted at in Section 2 and elaborated below, language does not simply activate prior sensorimotor experiences. Rather, it allows the creation of new task-relevant representations (see also [Deak, 2003](#)). So, imagine that you and a friend are looking for “Bill” in a crowd. “There he is!” your friend says. “I don’t see him” you respond. “He is wearing a red hat.” Looking for Bill initiates a search process that tunes the visual system to faces that comport with what you know about Bill’s appearance. The instruction that he is wearing a red hat dynamically retunes those features, further restricting the search. Indeed, simply hearing a word (over and above knowing what to search for) has been shown to improve visual search ([Lupyan, 2008b](#); [Lupyan & Spivey, 2010](#)).

4. Are words special?

Are words like any other cue in our environment? And if not, what makes them different? One way to begin answering this question is by comparing how words and other signals inform us about the current state of the world. What is the difference between some information being conveyed through language and information conveyed in other ways?

The question of how verbal and non-verbal cues differentially activate knowledge was addressed, albeit in a basic way, by Lupyan and Thompson-Schill (2012), who found that hearing words such as “dog” led to consistently faster recognition of subsequently presented pictures of the cued objects than was possible after hearing equally familiar non-verbal cues such as a dog bark. This “label-advantage” persisted for new categories of “alien musical instruments” for which participants learned to criterion either names or corresponding sounds—evidence that the advantage did not arise from the non-verbal cues being less familiar or inherently more difficult to process. Labels were not simply more effective at activating conceptual representations, but the representations activated by labels were systematically different from those activated by non-verbal sounds. Specifically, labels appeared to preferentially activate the most diagnostic features of the cued category as inferred by analyzing response patterns for different category exemplars. Thus, hearing “dog” appeared to selectively activate the features most central to recognizing dogs. As subsequent studies made clear ([Edmiston & Lupyan, 2015](#)), whereas non-verbal cues such as dog barks activated representations of specific dogs (and at specific times), while verbal labels appeared to activate more abstracted/prototypical representations. In short, labels, being categorical cues, are uniquely well-suited to activate categorical mental states.

There might appear to be an inconsistency in the idea that words both activate modal mental representations that are at the same time categorical. This apparent paradox can be clarified by work on canonical views of objects. Words activate perceptual representations of objects from particularly frequent ([Edelman & Bühlhoff, 1992](#); [Tarr, 1995](#); [Tarr & Pinker, 1989](#)) or informative perspectives ([Cutzu & Edelman, 1994](#)). Similarly, in the absence of contextual information to override it, people tend to generate visual representations (or simulations), from words that correspond to prototypical cases. These language-driven simulations can be thought of as schemas in that some degree of perceptual detail appears to be engaged (including shape, size, orientation, distance, and perhaps color), but finer details are typically not ([Holmes & Wolff, 2013](#); see [Bergen, 2012](#) for review). Thus, words activate internal states that are both prototypical and modally encoded.

The idea that language helps to move us from the particular to the categorical and abstract can be further elucidated with following example adapted from the philosopher Ruth Millikan:

There are many ways to recognize, say, rain. There is a way that rain feels when it falls on you, and a way that it looks through the window. There is a way that it sounds

falling on the rooftop, “retetetetetet,” and a way that it sounds falling on the ground, “shshshshsh.” And falling on English speakers, here is another way it can sound: “Hey, guys, it is raining!” (Millikan, 1998, p. 64)

Millikan’s insight is that language offers an additional way for the listener to identify the referent (in her words, “Recognizing a linguistic reference to a substance is just another way of re-identifying the substance itself,” Millikan, 1998, p. 64). But although there is, on our account, an overlap in the representations that are activated by the words and the perceptual experience, any experience of actual rain is an experience of *a particular instance* of rain. We can, of course, use language to be quite precise in specifying the kind of rain we mean. But we can also abstract away from almost all perceptual details and instantiate the kind of abstract rain that we can never experience directly: “Hey guys, it is raining.” Linguistic cues can thus abstract in a way that perceptual cues cannot. Language speaks in categories while perception speaks in particulars.

How does language come to have this property? Let us use the category of chairs as an example. Just as there is no rain in the abstract, there are no chairs in the abstract. There are highchairs and office-chairs, antique Victorian chairs, and uncomfortable modernist chairs. Each of our experiences with a chair is with a specific chair. Imagine now supplementing those sensorimotor experiences with a verbal label—“chair.” As “chair” is associated with various chairs, it becomes ever more strongly associated with properties most typical and diagnostic of chairs and *dissociated* from incidental features—those true of particular instances of chairs, but irrelevant to the category as a whole. A label like “chair” or “dog” can now activate a more categorical representation than is activated by seeing a chair or hearing a dog (Lupyan, 2008a, 2012, for a computational model).⁶ Using a label allows us to transcend—if temporarily—the concreteness of the real-world and to enter the world of abstractions.

To understand why abstraction is so important, think back to gathering coconuts. Every bit of the verbal instruction to climb a tree and gather some coconuts is an abstraction. How could such abstractions be communicated in the absence of language? Consider the difference between the linguistic expression and a pantomime of climbing, spinning, coconuts, etc. Not only do such pantomimes face the problem of distinguishing between similar-looking entities, but they necessarily depict particular types of climbing, a particular sized coconut, etc. In contrast, language allows the irrelevant details to be abstracted over (Lupyan, 2009). Our claim here is that language is not just a convenient way of achieving abstraction in communication, but rather that learning and using language is the primary means by which mental states are abstracted into such seemingly basic units as climb, tree, etc.

In addition to becoming a categorical cue, labels also become a highly effective cue for activating concepts in the absence of actual referents (indeed, arguably, most language refers to entities *not* currently in the speaker’s and listener’s environment). Just as we can use language to convey categorical information like “my favorite color is blue,” we can use language endogenously to activate and manipulate representations of spatially and temporally distant entities stripped away of task-irrelevant details. The more our

experience with a certain concept or domain is dominated by language use—both endogenous and exogenous—the more abstracted/categorical its representation may become.

5. Words, context, and compositionality

Word meanings (i.e., the mental states activated by words) are always interpreted within a context, both linguistic and non-linguistic. This means that it is difficult or even impossible to talk about invariant (“core”) meanings of words. For example, consider the word “banana.” One might imagine that the association between “banana” and the color yellow means that hearing “banana” activates, among other things, a perceptual representation of yellowness. And, in fact, reading the sentence “The bananas Mark bought looked ready to eat” appears to prime yellowness ([Connell & Lynott, 2009](#)). But it would be maladaptive for “banana” to activate yellowness in a reflexive way if only because bananas are not always yellow. And indeed, if a person instead reads that the bananas “were not ready to eat,” the priming effect disappears ([Connell & Lynott, 2009](#)). So, rather than viewing words as reflexively activating discrete bundles of features, it is more useful to think of words as inputs (cues) to a dynamic system that is always transitioning from one state to the next, the implication being that word meanings are intrinsically context-dependent ([Casasanto & Lupyan, 2015](#); [Elman, 2009](#)).

The role of context becomes clearer when we move from single words (“spin,” “coconut”) to larger utterances and engage the problem of compositionality. How do we get a person who knows about spinning and coconuts to put these meanings together and spin the coconut? This problem is partially solved by grammar. Just as words are cues to external referents, a grammar, as learned by each language-user, provides a recipe for combining words into larger meanings. Even though a person may have never heard a given combination of words before—“spin the coconut”—the phrase is similar in its form to other, previously heard phrases (it can be classified as a transitive construction, for example, [Goldberg, 2003](#)). Such computations act to construct the appropriate sensorimotor representation (sometimes called simulation). Clearly, structure matters. The simulations formed when processing “You handed Andy the pizza” and “Andy handed you the pizza” are measurably different even though they contain the same words ([Glenberg & Kaschak, 2002](#)). Combined, grammatical and pragmatic knowledge allow for flexible context-sensitive interpretation of words such that people can use language to program themselves or others to think and do things that are previously unknown, unthought, and unreal.

6. Implications for the evolution of language

We have argued that language is a powerful tool for programming the mind by helping to activate more abstract/categorical representations by affecting modal mental states. On this perspective, it becomes critical to view the evolution of language as not simply the

evolution of a communication system. After all, humans already *have* systems of acoustic and visual signals—postures, facial expressions, non-verbal vocalizations—that work just fine for other primates (Burling, 1993). The evolution of language, on this view, is seen as evolution of a *control* system for programming our own and other people’s minds.

The story, and it is just that, would go something like this: For reasons still poorly understood (but see Tomasello, 1999, 2008), humans diverged from other primates, becoming considerably more cooperative. This created a selective pressure on the emergence of a system to facilitate cooperation. Language is this system. As we have argued, language vastly simplifies teaching. By abstracting from task-irrelevant details (collapsing and discretizing the hypothesis space), language allows speakers and listeners to converge on task-relevant dimensions. Language promotes conceptual alignment by providing members of a community with a largely overlapping inventory of words promoting the learning of common categories. Via accumulation and transmission of knowledge across generations (facilitated by the linguistic encoding of continuous knowledge into a higher-fidelity categorical form), language promotes emergence of cultural institutions that depend on cooperation on ever larger scales (from tribe to villages and beyond).

There is an interesting corollary to this scenario. Although the evolutionary histories of other animals have evidently not led to the emergence of language, the control system aspects of language may be general enough to be used by other animals. There have been a number of attempts to teach aspects of language to bonobos (Savage-Rumbaugh, Shanker, & Taylor, 2001), a gray parrot (Pepperberg, 2013), and dogs (Griebel & Oller, 2012; Kaminski, Call, & Fischer, 2004; Pilley & Reid, 2011). It is often pointed out that these attempts were failures because they did not succeed in teaching other animals full-blown language. But what is much more interesting is that providing these animals with symbol-referent pairings absent from their usual environments appeared to augment their cognitive abilities in some of the same ways that language appears to augment the cognition of humans (e.g., Pepperberg & Carey, 2012; Thompson, Oden, & Boysen, 1997).

The finding that even gray parrots, animals rather distantly related to humans, can benefit from symbolic training suggests that many animals may have the cognitive capacities for learning aspects of natural languages (albeit simple ones). Still very little is known about the cognitive consequences of such learning. The absence of even simple languages among non-human animals may thus be more revealing about the particular circumstances in which language is adaptive (we put our bets on social factors such as cooperation) rather than special cognitive capacities unique to humans. Instead, the unique cognitive capacities of humans may emerge to a significant extent *due* to language and the cultural and educational institutions that language makes possible.

7. Conclusion

No training or programming will help human coconut-climbers leap from tree to tree like macaques. In contrast, while it takes but a couple words to invoke in a human the idea of collecting coconuts by twisting them off a tree, monkeys have to be trained step

by step using actual sensorimotor experience. Without the help of language-elicited abstractions, concepts like spin and tree, are difficult (perhaps impossible) to convey.

We have sketched out a view on which language is a critical tool for programming the mind. This idea is intellectually aligned with Clark's view of language as a "computational transformer which allows pattern-completing brains to tackle otherwise intractable classes of cognitive problems" (Clark, 1998, p. 163) and resonates with Dennett's idea of language as a "cognitive autostimulator" (Dennett, 1992, see also 1994).

Language, on this view, is not simply a communication system; it is a control system. With every word and larger construction we learn, we gain the ability to activate and reactivate mental content in systematic ways, allowing us to entertain concepts that could not cohere based on their perceptual features alone (e.g., EIGHTEEN) and to align conceptual representations with other members of a community even in the face of varied individual experiences. We can circumvent costly trial-and-error learning (Harlow, 1959) by following the instructions of others—an example of the benefit of "symbolic theft" over "sensorimotor toil" (Cangelosi, Greco, & Harnad, 2002). We can also tell *ourselves* what to do and simulate what it would be like to perform an action in different ways and then act with less uncertainty and less costly exploration of the solution space.

The evolutionary circumstances that led to the emergence of language may remain shrouded in mystery. But the idea of language as a tool for programming minds and the study of how language transforms mental representations is ripe for additional empirical investigation that will shed further light on the enduring puzzle of this human-defining trait.

Acknowledgments

The writing of this manuscript was supported by NSF-PAC grant 1331293 to GL and benefited greatly from conversations with the attendees of the 2012 Konrad Lorenz Institute Origins of Communication Workshop, especially from feedback by Bob McMurray.

Notes

1. One might argue that the crux of social learning is not programmability, but imitation. Is it not enough to just model a behavior and have the learner attempt to do it—repeatedly—until the desired outcome is met? It is not. Although the ability to imitate may be a pre-requisite for much of the social learning we do (Tomasello, 1999), learning through imitation is error prone even under the best of conditions. Consider, for example, the rather slow and fragile imitative learning of nut-cracking in chimpanzees (Marshall-Pescini & Whiten, 2008).
2. The power of language to guide learning of such procedural skills should not be underestimated. When one of us (GL) was learning to fly a plane, he did something he should not have and the instructor, alarmed, asked, "What were you thinking?"

GL responded that he thought his action would have a certain positive outcome, thus revealing to the instructor an incorrect theory. The instructor was then able to correct this theory by explaining why the outcome would be different than GL had expected. Language-guided learning of complex skills may remain difficult, but without language, the diagnosis and correction of incorrect beliefs (theories) would be far slower and more difficult (or perhaps impossible).

3. While it is possible to provide *some* perceptual or functional basis on which members of categories like DOG or KNIFE cohere together, it is not possible to provide such a basis for all the instances of the category of EIGHTEEN.
4. We can conceive all sorts of things for which we lack names or linguistic expressions. But these unnamed concepts are more likely to be variable across and within individuals, and to be less stable than concepts for entities we do name, as described below.
5. Also relevant to this discussion is the literature on similarity-based and rule-based learning that has emphasized the relationship between language and rule-learning (e.g., [Ashby & Maddox, 2005](#); [Minda & Miles, 2010](#)). All learning requires homing in on task-relevant dimensions. Rule-learning can be viewed as laying on one end of a continuum with purely similarity-based categories on one end, and formal rules on the other ([Pothos, 2005](#); [Sloutsky, 2010](#); [Lupyan et al., 2012](#); [Perry & Lupyan, 2014](#)).
6. What about other unambiguous information like a dog bark as a cue to dogs? The critical difference between “dog” and a barking sound appears to lie in *motivation*. A dog’s bark varies in a lawful way with perceptual properties of the dog, for example, larger dogs produce lower pitched barks. The word “dog” also varies but unlike the bark, the exact form of the word is not motivated by perceptual details of the referent. We can tell a lot about the speaker by the way he or she says “dog” but not about the dog. This makes non-verbal cues like barking highly effective indices, but bad referential cues ([Edmiston & Lupyan, 2013](#); under review).

References

- [Addyman, C., & Mareschal, D. \(2010\). The perceptual origins of the abstract same/different concept in human infants. *Animal Cognition*, 13\(6\), 817–833. doi: 10.1007/s10071-010-0330-0.](#)
- [Althaus, N., & Mareschal, D. \(2014\). Labels direct infants’ attention to commonalities during novel category learning. *PLoS ONE*, 9\(7\), e99670. doi: 10.1371/journal.pone.0099670.](#)
- [Amsel, B. D., Urbach, T. P., & Kutas, M. \(2014\). Empirically grounding grounded cognition: The case of color. *NeuroImage*, 99C, 149–157. doi: 10.1016/j.neuroimage.2014.05.025.](#)
- [Ashby, F. G., & Maddox, W. T. \(2005\). Human category learning. *Annual Review of Psychology*, 56\(1\), 149–178. doi: 10.1146/annurev.psych.56.091103.070217.](#)
- [Baddeley, A. D., Chincotta, D., & Adlam, A. \(2001\). Working memory and the control of action: Evidence from task switching. *Journal of Experimental Psychology: General*, 130\(4\), 641–657. doi:10.1037/0096-3445.130.4.641.](#)
- [Barsalou, L. W. \(1999\). Perceptual symbol systems. *The Behavioral and Brain Sciences*, 22\(4\), 577–609; discussion 610–660.](#)

- Bergen, B. K. (2012). *Louder than words: The new science of how the mind makes meaning*. New York: Basic Books.
- Bloom, P. (2001). Roots of word learning. In M. Bowerman & S. C. Levinson (Eds.), *Language acquisition and conceptual development* (pp. 159–181). Cambridge, UK: Cambridge University Press.
- Bloom, P., & Keil, F. C. (2001). Thinking through language. *Mind & Language*, *16*(4), 351–367.
- Borges, J. L. (1999). *Collected fictions*. New York: Penguin (Non-Classics).
- Burling, R. (1993). Primate calls, human language, and nonverbal communication. *Current Anthropology*, *34*(1), 53, 25.
- Cangelosi, A., Greco, A., & Harnad, S. (2002). Symbol grounding and the symbolic theft hypothesis. In A. Cangeosi and M. Parisi (Eds.), *Simulating the evolution of language* (pp. 191–210). London: Springer.
- Casasanto, D., & Lupyan, G. (2014). All concepts are ad hoc concepts. In E. Margolis & S. Laurence (Eds.), *Concepts: New directions*. Cambridge, MA: MIT Press.
- Casasola, M. (2005). Can language do the driving? The effect of linguistic input on infants' categorization of support spatial relations. *Developmental Psychology*, *41*(1), 183–192. doi: 10.1037/0012-1649.41.1.188.
- Cassirer, E. (1962). *An essay on man: An introduction to a philosophy of human culture*. New Haven, CT: Yale University Press.
- Castro, L., & Wasserman, E. A. (2013). Humans deploy diverse strategies in learning same-different discrimination tasks. *Behavioural Processes*, *93*, 125–139. doi: 10.1016/j.beproc.2012.09.015.
- Clark, A. (1998). Magic words: How language augments human computation. In P. Carruthers & J. Boucher (Eds.), *Language and thought: Interdisciplinary themes* (pp. 162–183). New York: Cambridge University Press.
- Collins, J. A., & Curby, K. M. (2013). Conceptual knowledge attenuates viewpoint dependency in visual object recognition. *Visual Cognition*, *21*(8), 945–960. doi: 10.1080/13506285.2013.836138.
- Connell, L., & Lynott, D. (2009). Is a bear white in the woods? Parallel representation of implied object color during language comprehension. *Psychonomic Bulletin & Review*, *16*(3), 573–577. doi: 10.3758/PBR.16.3.573.
- Cutzu, F., & Edelman, S. (1994). Canonical views in object representation and recognition. *Vision Research*, *34*(22), 3037–3056. doi: 10.1016/0042-6989(94)90277-1.
- Deak, G. O. (2003). The development of cognitive flexibility and language abilities. *Advances in child development and behavior* (pp. 271–327). San Diego: Academic Press.
- Dennett, D. (1992). *Consciousness explained* (1st ed.). Boston: Back Bay Books.
- Dennett, D. (1994). The role of language in intelligence. In *What is intelligence? The Darwin college lectures* (pp. 161–178). Cambridge, England: Cambridge University Press.
- Edelman, S., & Bühlhoff, H. H. (1992). Orientation dependence in the recognition of familiar and novel views of three-dimensional objects. *Vision Research*, *32*(12), 2385–2400. doi: 10.1016/0042-6989(92)90102-O.
- Edmiston, P., & Lupyan, G. (2015). What makes words special? Words as unmotivated cues. *Cognition*, *143*, 93–100.
- Elman, J. L. (2009). On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science*, *33*(4), 547–582. doi: 10.1111/j.1551-6709.2009.01023.x.
- Emerson, M. J., & Miyake, A. (2003). The role of inner speech in task switching: A dual-task investigation. *Journal of Memory and Language*, *48*(1), 148–168.
- Estes, Z., Verges, M., & Barsalou, L. W. (2008). Head up, foot down: Object words orient attention to the objects' typical location. *Psychological Science*, *19*(2), 93–97. doi: 10.1111/j.1467-9280.2008.02051.x.
- Evans, V. (2009). *How words mean lexical concepts, cognitive models, and meaning construction*. Oxford, England: Oxford University Press.
- Farrell, B. (1985). Same different judgments: A review of current controversies in perceptual comparisons. *Psychological Bulletin*, *98*(3), 419–456.
- Fodor, J. A. (2010). *LOT 2: The language of thought revisited*. Oxford, England: Oxford University Press.
- Frank, M. C., Everett, D. L., Fedorenko, E., & Gibson, E. (2008). Number as a cognitive technology: Evidence from Pirahã language and cognition. *Cognition*, *108*(3), 819–824. doi: 10.1016/j.cognition.2008.04.007.
- Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, *9*(3), 558–565.

- Glenberg, A. M., Lopez-Mobilia, G., McBeath, M., Toma, M., Sato, M., & Cattaneo, L. (2010). Knowing beans: Human mirror mechanisms revealed through motor adaptation. *Frontiers in Human Neuroscience*, 4. doi: 10.3389/fnhum.2010.00206.
- Goldberg, A. E. (2003). Constructions: A new theoretical approach to language. *Trends in Cognitive Sciences*, 7(5), 219–224.
- Gomila, T. (2011). *Verbal minds: Language and the architecture of cognition* (1 edition.). Amsterdam: Elsevier.
- Gordon, P. (2004). Numerical cognition without words: Evidence from Amazonia. *Science*, 306(5695), 496–499.
- Griebel, U., & Oller, D. K. (2012). Vocabulary learning in a Yorkshire terrier: Slow mapping of spoken words. *PLoS ONE*, 7(2), e30182. doi: 10.1371/journal.pone.0030182.
- Harlow, H. (1959). Learning set and error factor theory. *Psychology: A Study of a Science*, 2, 492–537.
- Harnad, S. (2005). Cognition is categorization. In H. Cohen, & C. Lefebvre (Eds.), *Handbook of categorization in cognitive science* (pp. 20–45). San Diego, CA: Elsevier Science.
- Hays, P. R. (2000). From the Jurassic dark: Linguistic relativity as evolutionary necessity. In M. Pütz & M. Verspoor (Eds.), *Explorations in linguistic relativity* (pp. 173–196). Philadelphia: John Benjamins.
- Holmes, K. J., & Wolff, P. (2013). Spatial language and the psychological reality of schematization. *Cognitive Processing*, 14(2), 205–208. doi: 10.1007/s10339-013-0545-5.
- Jame, W. (1890). *Principles of psychology*. Vol. 1. New York: Holt.
- Kaminski, J., Call, J., & Fischer, J. (2004). Word learning in a domestic dog: Evidence for “fast mapping.” *Science*, 304(5677), 1682–1683. doi: 10.1126/science.1097859.
- Karbach, J., Kray, J., & Hommel, B. (2011). Action-effect learning in early childhood: Does language matter? *Psychological Research*, 75(4), 334–340. doi: 10.1007/s00426-010-0308-1.
- Karmiloff-Smith, A. (1990). Constraints on representational change: Evidence from children’s drawing. *Cognition*, 34(1), 57–83.
- Kiefer, M., & Pulvermüller, F. (2012). Conceptual representations in mind and brain: Theoretical developments, current evidence and future directions. *Cortex*, 48(7), 805–825. doi: 10.1016/j.cortex.2011.04.006.
- Kosslyn, S. M., Ganis, G., & Thompson, W. L. (2001). Neural foundations of imagery. *Nature Reviews Neuroscience*, 2(9), 635–642.
- Levinson, S. C. (1997). From outer to inner space: Linguistic categories and non-linguistic thinking. In J. Nuyts & E. Pederson (Eds.), *Language and conceptualization* (pp. 13–45). Cambridge: Cambridge University Press.
- Li, P., & Gleitman, L. (2002). Turning the tables: Language and spatial reasoning. *Cognition*, 83(3), 265–294.
- Lupyan, G. (2008a). From chair to “chair”: A representational shift account of object labeling effects on memory. *Journal of Experimental Psychology: General*, 137(2), 348–369.
- Lupyan, G. (2008b). The conceptual grouping effect: Categories matter (and named categories matter more). *Cognition*, 108(2), 566–577.
- Lupyan, G. (2009). Extracommunicative functions of language: Verbal interference causes selective categorization impairments. *Psychonomic Bulletin & Review*, 16(4), 711–718. doi: 10.3758/PBR.16.4.711.
- Lupyan, G. (2012). What do words do? Towards a theory of language-augmented thought. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 57, pp. 255–297). Waltham, MA: Academic Press. Available at: <http://www.sciencedirect.com/science/article/pii/B9780123942937000078>
- Lupyan, G., & Casasanto, D. (2014). Meaningless words promote meaningful categorization. *Language and Cognition, FirstView*, 1–27. doi: 10.1017/langcog.2014.21.
- Lupyan, G., & Mirman, D. (2013). Linking language and categorization: Evidence from aphasia. *Cortex*, 49(5), 1187–1194. doi: 10.1016/j.cortex.2012.06.006.
- Lupyan, G., Mirman, D., Hamilton, R. H., & Thompson-Schill, S. L. (2012). Categorization is modulated by transcranial direct current stimulation over left prefrontal cortex. *Cognition*, 124(1), 36–49. doi: 10.1016/j.cognition.2012.04.002.

- Lupyan, G., Rakison, D. H., & McClelland, J. L. (2007). Language is not just for talking: Labels facilitate learning of novel categories. *Psychological Science, 18*(12), 1077–1082.
- Lupyan, G., & Spivey, M. J. (2010). Redundant spoken labels facilitate perception of multiple items. *Attention, Perception, & Psychophysics, 72*(8), 2236–2253. doi: 10.3758/APP.72.8.2236.
- Lupyan, G., & Swingle, D. (2012). Self-directed speech affects visual search performance. *Quarterly Journal of Experimental Psychology, 65*(6), 1068–1085. doi: 10.1080/17470218.2011.647039.
- Lupyan, G., & Thompson-Schill, S. L. (2012). The evocative power of words: Activation of concepts by verbal and nonverbal means. *Journal of Experimental Psychology. General, 141*(1), 170–186. doi: 10.1037/a0024904.
- Lupyan, G., & Ward, E. J. (2013). Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences, 110*(35), 14196–14201. doi: 10.1073/pnas.1303312110.
- Mahon, B. Z., & Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology. Paris, 102*(1–3), 59–70. doi: 10.1016/j.jphysparis.2008.03.004.
- Malt, B. C., Gennari, S. P., Imai, M., Ameel, E., Saji, N., & Majid, A. (2015). Where are the concepts? What words can and can't reveal. In E. Margolis & S. Laurence (Eds.), *Concepts: New directions*. Cambridge: MIT Press.
- Mandler, J. (2004). *The foundations of mind? Origins of conceptual thought: Origins of conceptual thought*. New York: Oxford University Press.
- Marmor, G. S. (1978). Age at onset of blindness and the development of the semantics of color names. *Journal of Experimental Child Psychology, 25*(2), 267–278. doi: 10.1016/0022-0965(78)90082-6.
- Marshall-Pescini, S., & Whiten, A. (2008). Social learning of nut-cracking behavior in East African sanctuary-living chimpanzees (*Pan troglodytes schweinfurthii*). *Journal of Comparative Psychology, 122*(2), 186–194. doi: 10.1037/0735-7036.122.2.186.
- McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review, 119*(4), 831–877. doi: 10.1037/a0029872.
- Millikan, R. G. (1998). A common structure for concepts of individuals, stuffs, and real kinds: More mama, more milk, and more mouse. *Behavioral and Brain Sciences, 21*(1), 55–65.
- Minda, J. P., & Miles, S. J. (2010). The influence of verbal and nonverbal processing on category learning. In B. H. Ross (Ed.), *Psychology of learning and motivation* (Vol. 52, pp. 117–162). Waltham, MA: Academic Press. Available at: <http://www.sciencedirect.com/science/article/pii/S0079742110520036>
- Nazzi, T., & Gopnik, A. (2001). Linguistic and cognitive abilities in infancy: When does language become a tool for categorization? *Cognition, 80*(3), B11–B20.
- Ozçalışkan, S., Goldin-Meadow, S., Gentner, D., & Mylander, C. (2009). Does language about similarity play a role in fostering similarity comparison in children? *Cognition, 112*(2), 217–228. doi: 10.1016/j.cognition.2009.05.010.
- Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences, 31*(02), 109–130. doi: 10.1017/S0140525X08003543.
- Pepperberg, I. M. (2013). Abstract concepts: Data from a grey parrot. *Behavioural Processes, 93*, 82–90. doi: 10.1016/j.beproc.2012.09.016.
- Pepperberg, I. M., & Carey, S. (2012). Grey parrot number acquisition: The inference of cardinal value from ordinal position on the numeral list. *Cognition, 125*(2), 219–232. doi: 10.1016/j.cognition.2012.07.003.
- Perry, L. K., & Lupyan, G. (2014). The role of language in multi-dimensional categorization: Evidence from transcranial direct current stimulation and exposure to verbal labels. *Brain and Language, 135*, 66–72. doi: 10.1016/j.bandl.2014.05.005.
- Pilly, J. W., & Reid, A. K. (2011). Border collie comprehends object names as verbal referents. *Behavioural Processes, 86*(2), 184–195. doi: 10.1016/j.beproc.2010.11.007.

- Plunkett, K., Hu, J.-F., & Cohen, L. B. (2008). Labels can override perceptual categories in early infancy. *Cognition*, *106*(2), 665–681. doi: S0010-0277(07)00108-4.
- Pothos, E. (2005). The rules versus similarity distinction. *Behavioral and Brain Sciences*, *28*, 1–49.
- Pulvermüller, F., Shtyrov, Y., & Hauk, O. (2009). Understanding in an instant: Neurophysiological evidence for mechanistic language circuits in the brain. *Brain and Language*, *110*(2), 81–94. doi: 10.1016/j.bandl.2008.12.001.
- Pulvermüller, F., Shtyrov, Y., & Ilmoniemi, R. (2005). Brain signatures of meaning access in action word recognition. *Journal of Cognitive Neuroscience*, *17*(6), 884–892. doi: 10.1162/0898929054021111.
- Robinson, C. W., Best, C. A., Deng, W. S., & Sloutsky, V. M. (2012). The role of words in cognitive tasks: What, when, and how? *Frontiers in Psychology*, *3*. doi: 10.3389/fpsyg.2012.00095.
- Rumelhart, D. E., Smolensky, D., McClelland, J. L., & Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In J. L. McClelland, D. E. Rumelhart, & PDP Research Group (Eds.), *Parallel distributed processing vol II* (pp. 7–57). Cambridge, MA: MIT Press.
- Savage-Rumbaugh, S., Shanker, S. G., & Taylor, T. J. (2001). *Apes, language, and the human mind*. Oxford, England: Oxford University Press.
- Sloutsky, V. M. (2010). From perceptual categories to concepts: What develops? *Cognitive Science*, *34*(7), 1244–1286. doi: 10.1111/j.1551-6709.2010.01129.x.
- Snedeker, J., & Gleitman, L. (2004). Why is it hard to label our concepts? In D. G. Hall & S. R. Waxman (Eds.), *Weaving a lexicon* (illustrated ed.) (pp. 257–294). Cambridge, MA: MIT Press.
- Tarr, M. J. (1995). Rotating objects to recognize them: A case study on the role of viewpoint dependency in the recognition of three-dimensional objects. *Psychonomic Bulletin & Review*, *2*(1), 55–82. doi: 10.3758/BF03214412.
- Tarr, M. J., & Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, *21*(2), 233–282. doi: 10.1016/0010-0285(89)90009-1.
- Taylor, L. J., & Zwaan, R. A. (2008). Motor resonance and linguistic focus. *Quarterly Journal of Experimental Psychology*, *61*(6), 896–904. doi: 10.1080/17470210701625519.
- Thompson, R. K., Oden, D. L., & Boysen, S. T. (1997). Language-naïve chimpanzees (Pan troglodytes) judge relations between relations in a conceptual matching-to-sample task. *Journal of Experimental Psychology. Animal Behavior Processes*, *23*(1), 31–43.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
- Tomasello, M. (2008). *Origins of human communication*. Cambridge, MA: MIT Press.
- Vygotsky, L. (1962). *Thought and language*. Cambridge, MA: MIT Press.
- Wagner, K., Dobkins, K., & Barner, D. (2013). Slow mapping: Color word learning as a gradual inductive process. *Cognition*, *127*(3), 307–317. doi: 10.1016/j.cognition.2013.01.010.
- Walker, C. M., & Gopnik, A. (2014). Toddlers infer higher-order relational principles in causal learning. *Psychological Science*, *25*(1), 161–169. doi: 10.1177/0956797613502983.
- Waxman, S. R., & Markow, D. B. (1995). Words as invitations to form categories: Evidence from 12- to 13-month-old infants. *Cognitive Psychology*, *29*(3), 257–302.
- Whorf, B. L. (1956). *Language, thought, and reality*. Cambridge, MA: MIT Press.
- Winter, B., & Bergen, B. (2012). Language comprehenders represent object distance both visually and auditorily. *Language and Cognition*, *4*(01), 1–16. <http://doi.org/10.1515/langcog-2012-0001>
- Yee, E., Chrysikou, E. G., Hoffman, E., & Thompson-Schill, S. L. (2013). Manual experience shapes object representations. *Psychological Science*, *24*(6), 909–919. doi: 10.1177/0956797612464658.
- Young, M. E., & Wasserman, E. A. (2001). Entropy and variability discrimination. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *27*(1), 278–293.
- Zwaan, R. A., Madden, C. J., Yaxley, R. H., & Aveyard, M. E. (2004). Moving words: Dynamic representations in language comprehension. *Cognitive Science*, *28*(4), 611–619. doi: 10.1207/s15516709cog2804_5.
- Zwaan, R. A., Stanfield, R. A., & Yaxley, R. H. (2002). Language comprehenders mentally represent the shapes of objects. *Psychological Science*, *13*(2), 168–171. doi: 10.1111/1467-9280.00430.