19   All Concepts Are Ad Hoc Concepts

Daniel Casasanto and Gary Lupyan

19.1   Overview

To explain how people think and communicate, cognitive scientists posit a repository of concepts, categories, and word meanings that are stable across time and shared across individuals. But if concepts are stable, how can people use them so flexibly? Here we explore a possible answer: maybe this stability is an illusion. Perhaps all concepts, categories, and word meanings (CC&Ms) are constructed ad hoc, each time we use them. On this proposal, which we call the *ad hoc cognition* (AHC) framework, all words are infinitely polysemous, all communication is "good enough," and no idea is ever the same twice. The details of people's ad hoc CC&Ms are determined by the way retrieval cues (such as words) interact with the physical, social, and linguistic context. Commonalities across instantiations of CC&Ms yield some emergent stability and create the illusion of context-independent core properties. Here we argue that even the most stable-seeming CC&Ms are instantiated via the same processes as those that are more obviously ad hoc, and they vary (a) from one microsecond to the next within a given instantiation, (b) from one instantiation to the next within an individual, and (c) from person to person and group to group as a function of people's experiential history. If this is true, then a central goal of research on language and cognition should be to elucidate the fleeting, idiosyncratic neurocognitive representations that people actually use for thinking and communicating, rather than to discern the nature and origin of context-independent CC&Ms, which, we argue, only exist as theoretical abstractions. Thinking depends on brains, and brains are always changing; therefore thoughts are always changing. Rather than trying to explain concepts, categories, and word meanings as things that we *have in* our minds, like entries in a mental dictionary or mental encyclopedia, it may be more fruitful to build theories of conceptualizing, categorizing, and constructing word meanings: things that we *do with* our minds.

El

19.2   Concepts, Categories, and Word Meanings as Scientific Constructs

What's in our minds? How is our knowledge organized, stored, and used? Two millennia ago, these questions could only be addressed through allegory. For example, Plato considered whether the mind might be like an aviary, and each bird in it a piece of knowledge (*Theaetetus*, ca. 360 BCE). When the birds are not in use, they are free to flap about the cage of our memories. To use a piece of knowledge, we just need to catch (i.e., retrieve) the right bird.

It was clear then, as it is now, that there are not actually birds in our heads. But today, many people believe that the mind is populated by much more elusive creatures: concepts, categories, and word meanings. Aspects of the bird-catching model of the mind remain strongly entrenched insomuch as CC&Ms are often characterized as discrete entities that exist fully formed, even when we are not using them, to be summoned as needed.

According to the AHC framework we sketch here, concepts, categories, and word meanings only exist as "analytic fictions created by those who study them" (Barsalou 1987, 119), rooted in folk-psychological notions (e.g., concepts are entries in a mental encyclopedia, word meanings in a mental dictionary). As scientific constructs, these fictions can be useful for describing the neurocognitive states and processes that cognitive scientists seek to characterize, as long as one's "concept of X" or "the meaning of word Y" are not mistaken for natural kinds whose true essence can be discovered. Barrett, Mesquita, and Smith (2010) observe:

Scientific disciplines categorize. They divide their universe of interests into groupings of "kinds," name them, and then set about the business of understanding those kinds. … This categorization process functions like a sculptor's chisel, dividing up the world into figure and ground, leading scientists to attend to certain features and to ignore others. One consequence of scientific categorization is that we sometimes *essentialize* our subject matter, then search for evidence of those essences, without considering how context might influence or contribute to its very nature. (1; italics added)

Many researchers in the physical and biological sciences have shed the illusion that there exist observer-independent natural kinds. This belief is untenable in light of their fields' emerging understanding of ways in which the act of observing the world changes it, and in light of repeated illustrations that today's scientific constructs may become tomorrow's quaint footnotes in history. "Natural kinds" like phlogiston (the fire-like element chemists once believed to cause combustion reactions), the élan vital (the animating force believed, as recently as the twentieth century, to separate the living from the nonliving), and the brontosaurus (a chimera accidentally created out of two other species' bones) were once the subjects of serious scientific inquiry and popular belief. For many researchers, the scientist's job is now to construct useful models of the

El

natural world, rather than to discern some latent "true" structure in it (see Chomsky 2002; Dunham and Banaji 2010).

We suggest that progress toward a scientific understanding of thinking and language use would be accelerated if researchers were to consider that (a) CC&Ms are human creations rather than natural kinds, and (b) the patterns of neurocognitive activity that CC&Ms are meant to provide a model of are constructed ad hoc and shaped by the physical and social contexts in which they are instantiated, on various time scales. Aspects of this proposal overlap with other proposals that emphasize the importance of context and the dynamism of language and thought (e.g., Barsalou 1987; Churchland 1986; Clark 1996, 1997; Elman 2004, 2009; Evans 2009; Hampton 2012; Machery 2009; Prinz 2002; Rogers and McClelland 2004; Smith and Samuelson 1997; Spivey 2007; Taylor and Zwaan 2009; Weiskopf 2009; Wittgenstein 1953).

We argue that it is necessary not only to acknowledge that thinking happens on the fly, but also to abandon the idea that CC&Ms have stable *cores* or *defaults* that people simply access whenever they instantiate CC&Ms (cf. Armstrong, Gleitman, and Gleitman 1983; Osherson and Smith 1981). Instead, each instantiation of CC&Ms is constructed on the basis of retrieval cues (internal and external) embedded in an ever-changing physical, social, and internal biological context.

Proposals in various literatures have argued that some circumstances can lead to the creation of ad hoc concepts (e.g., Allott and Textor 2012) or categories (e.g., Barsalou 1983). According to these proposals, however, ad hoc concepts and categories are created only when "common" concepts and categories—from which the ad hoc versions are qualitatively different—are unsuitable. We will argue here that common CC&Ms *are* ad hoc CC&Ms.

Across the cognitive sciences, concepts, categories, and word meanings are treated as distinct constructs for some purposes but as interdependent constructs for others: typically, *concepts* are what allow us to categorize, *categories* are the extensions of concepts, and *words have meanings* by virtue of activating (lexical) concepts. Accordingly, we will treat CC&Ms as indistinguishable throughout much of the chapter (thus their amalgamation in the term "CC&M"), but we will also treat them separately in order to engage with the distinct literatures about concepts, categories, and word meanings.[1]

In what follows, we first sketch an account of CC&Ms according to the ad hoc cognition framework, and then review evidence of the pervasive instability of mental representations that motivates this view. We illustrate this instability on three overlapping

1. We recognize that notions like concepts and categories, and the desiderata for what work they should do, may differ between psychologists and philosophers (Hampton 2012; Margolis and Laurence 2008; Machery 2009). In general, we will be writing from the perspective of researchers primarily interested in advancing theories in psychology and cognitive neuroscience.

El

time scales: *activation dynamics*, *local context effects*, and *experiential relativity.* We then offer some explanations of why, in spite of this demonstrable variability, people find it hard to escape the intuition that invariant CC&Ms exist, and briefly suggest ways in which the focus of research in the cognitive sciences should change in light of the ad hoc nature of cognition.

19.2.1   From *Concepts* to *Conceptualizing*

We will use the term *concept* to mean a dynamic pattern of information that is made active in memory transiently, as needed, in response to internally generated or external cues. Activating a "concept of X" is thinking about X; this should be uncontroversial. More controversially perhaps, we suggest that concepts exist only when they are being used. Concepts are not something we *have in* the mind, they are something we *do with* the mind. The word *concept* is, itself, problematic (though hard to avoid) insomuch as naming something with a noun seems to imply it is an object, but *conceptualizing* is a process (see Barsalou, Wilson, and Hasenkamp 2010; Spivey 2007).

Rather than a process of accessing a preformed package of knowledge, instantiating a concept is always a process of activating an ad hoc network of stored information in response to cues in context. These cues may be verbal (e.g., the word *dog*) or nonverbal (e.g., a barking sound), exogenous (coming from the outside world) or endogenous (internally generated). Cues may be processed consciously or unconsciously. Even if the exact same cue could be experienced repeatedly, the context in which it is experienced can never be re-created exactly, since one's internal neurocognitive state is part of the context. At minimum, the contexts at time 1 and time 2 differ because, at time 2, one's neurocognitive state has been altered by the processes that occurred at time 1.

As such, no two instantiations of "the same" concept are ever identical. In William James's words, "a permanently existing 'idea' … which makes its appearance before the footlights of consciousness at periodic intervals is as mythological an entity as the Jack of Spades" (1890, 230). Different responses to the same cue (or similar cues) may be very similar, and this similarity may contribute to the illusion that concepts are stable over time, but as we argue in section 19.5.3, the unity of concepts has been greatly overestimated.

19.2.2   From *Categories* to *Categorizing*

In his landmark paper introducing the term *ad hoc categories* to psychology, Barsalou (1983) showed that goal-derived, for-the-nonce categories like *things to sell at a garage sale* are structured similarly to "common" categories like *vehicles*. According to Barsalou, the main distinction between ad hoc and common categories is that only common categories are well entrenched in memory. As evidence for this distinction, he reported single dissociations between behavioral responses to stimuli designed to activate ad hoc versus common categories. Arguably, however, the apparently categorical

differences these studies showed could be artifacts of using categorical experimental designs, which artificially dichotomize a continuum of entrenchment in memory. Consistent with this assertion, Barsalou wrote that raising the frequency with which an ad hoc category is activated can turn it into a common category (at least for the individual who's doing the categorizing), and lowering the frequency of use can turn a common category into an ad hoc category.

Since there appears to be no categorical difference between these supposedly different types, we suggest that the terminological distinction between common (i.e., stable) and ad hoc categories has outlived its usefulness. In later writing, Barsalou (1987) appears amenable to the proposal that all categorizing is ad hoc:

> Knowledge in long-term memory … is relatively undifferentiated and continuous. … It may be extremely difficult, if not impossible, to identify where the knowledge for a particular category in long-term memory begins and ends. To the extent that this is true, it is hard to imagine how there could be invariant representations for categories stored in long-term memory. (121)

On our view, *categorizing* (i.e., forming equivalence classes that support generalization and inference) is surely among the most important functions of the brain and mind, but *categories* as static structures in long-term memory—or as observer-independent collections of entities in the world—are a theoretical dead end.

19.2.3 From Accessing Word Meanings to Constructing Them

On standard views of language understanding, words have a finite variety of meanings and senses, and the context determines which of these gets "accessed." Accessing a word's meaning (or sense) means retrieving the appropriate concept (e.g., Jackendoff 2002; but see Elman 2004; Evans 2009). From the concept, activation often spreads to other information in long-term memory, which on many theories, is held to be *associated* with the concept or word meaning but is not *part of* it (e.g., Mahon and Caramazza 2008).

By contrast, within the AHC framework, word forms are cues to activate stored information, as needed and as determined by the specifics of the internal neurocognitive and external physical and social contexts. There is no principled distinction between information cued by a word-in-context that is *constitutive* of the meaning versus merely *associated* with it. The pattern of activity that constitutes a word's ad hoc meaning can be called the ad hoc concept that it activates in a given instance, but there is no memorized mapping between word forms and concepts. AHC, therefore, rejects an assumption that cuts across many otherwise divergent theories of language: that word meanings consist of memorized, conventionalized form-meaning pairings (e.g., Evans 2009; Jackendoff 2002; Langacker 2008; Snedeker and Gleitman 2004).

In short, according to AHC, words do not *have* meanings; rather, a word-in-context is a cue to construct what can be called its meaning for a given instantiation

El

(see also Elman 2004). The meaning that gets activated in response to a word depends on the individual's history of processing this cue, as well as its current context. Meanings may be similar across instances, but they are always constructed in response to words-in-context functioning as complex retrieval cues. Since a given internal and external context can never be duplicated, no word ever has exactly the same meaning twice.

There is no principled distinction between the role that a word plays in cuing a neurocognitive representation and the role that its context plays. The impossibility of such a distinction becomes clear as soon as language use is considered in more ecologically valid settings than laboratory experiments involving exposure to single words. People don't activate meanings in response to words: they activate meanings in response to words-in-context. (This is true even when words appear in isolation, in the lab, as we will argue in the "local context" section, 19.5.2.)

Because words are not normally isolated but occur in the context of other linguistic, physical, and social (e.g., pragmatic) cues, people do not need to wait to hear a word in order to start understanding it: language users *predict* meaning. In a seminal study, participants who were presented with an array of objects (e.g., pictures of a ball, a train, a truck, and a cake) were likely to start looking at the cake upon hearing "*the boy will eat …,*" apparently finishing the sentence in their minds before hearing "*the cake*" (Altmann and Kamide 1999). The proportion of looks to the cake increased smoothly over the time that participants heard *eat … the … cake,* indicating that their representations of cake were constructed incrementally (see also DeLong, Urbach, and Kutas 2005; Dikker et al. 2010). Although the context rarely fully constrains what will follow, ordinary language understanding relies on the same predictive processes shown in this study: people begin to partially construct the meanings of words before they hear (or read) them, tuning the network of activated information gradually as the context constrains their predictions.

19.3   There Are No Context-Independent CC&Ms

Is a word's meaning *always* co-determined by its context, or are some words' meanings (or some aspects of word meaning) invariant across all contexts? Barsalou (1982) proposed that some features of concepts and word meanings are *context independent* (CI) and are activated maximally and obligatorily whenever a word is perceived, irrespective of its context. Other features are *context dependent* (CD): activated by the context, and *not* by the word that names the concept. The notion of CI features follows from the standard assumption that words activate fixed concepts, or conceptual cores, and the notion of CD features follows from the assumption that words are separable from their contexts; neither assumption is compatible with AHC. Even while arguing for the

flexibility of concepts, Barsalou (1987) maintained the distinction between CD and CI properties, which still figures prominently in theories of concepts decades later (e.g., Machery 2009).

The experimental results supporting the CD-CI distinction (Barsalou 1982), however, like the results supporting the distinction between ad hoc and common categories, may primarily reflect the categorical distinctions imposed on the stimulus materials by the experimenters: if you ask a dichotomous question, often you get a dichotomous answer.

The properties that Barsalou identified as CD and CI appear to lie along a continuum of automaticity and context dependence—but it seems unlikely that any are context invariant. Barsalou suggested, for example, that the property "has a smell" is a context-invariant property of skunks, which is always maximally activated in response to this word. Is this true? In Barsalou's (1982) experiments, the target word (e.g., skunk) was always presented *before* the contextual information that determined whether the putatively invariant property was relevant. What if the contextual information had come first? If the prior context had made it clear that the skunk was a toy skunk, or a taxidermied skunk, or a lithograph of a skunk, wouldn't smell-related information have been activated less strongly than in other contexts involving, say, lifted tails? Crucially, Barsalou (1982) compared the activation of putatively context-invariant properties in contexts that supported the property versus neutral contexts: there were no tests of whether the supposedly CI properties were equally activated in contexts that were clearly *unsupportive* of the property, which would provide a stronger test of the property's context independence.

We suggest that there is no principled way to distinguish between context-dependent and context-independent information cued by words because there is no such thing as context-independent information. A given word may tend to activate some information more frequently and more automatically than other information, but all information is context dependent. Consistent with this assertion, Barsalou (1982, 1987) suggested that, over the course of experience, properties that start out as *context independent* in an individual's mind can become context dependent, and vice versa, belying any categorical distinction.[2]

2. In personal communication with D. C. (October 20, 2012), L. Barsalou agreed that the CI-CD distinction has outlived its usefulness. We consider Barsalou's studies establishing the constructs of ad hoc categories (1983) and context-dependent features of concepts (1982) to be of great theoretical importance, even though we disagree with their conclusions. These forward-thinking studies arguing that *some* CC&Ms are ad hoc paved the way for the proposal that *all* CC&Ms are ad hoc. Barsalou's contemporary writings on the dynamism of concepts (e.g., Barsalou, Wilson, and Hasenkamp 2010) are compatible with what we propose here and provide inspiration to further develop the AHC framework.

Demonstrations of context dependence are abundant (see section 19.5); here we focus on Barsalou's (1982) study because it represents a considered attempt to demonstrate that some properties of concepts are truly context independent, as posited by some influential theories of word meaning (e.g., Katz and Postal 1964), forming a fixed core of a word's meaning that is activated "on all occasions" and is "unaffected by contextual relevance" (Barsalou 1982, 82). In light of subsequent research, however, it seems unlikely that any aspects of concepts or word meanings fulfill these criteria.

The Stroop effect (Stroop 1935), which has been replicated in more than five hundred experiments (Besner, Stolz, and Boutilier 1997), is often interpreted as evidence that concepts are activated automatically in response to words. In a classic Stroop task, participants see the names of colors printed in ink that is either congruent in color (e.g., *blue* printed in blue ink) or incongruent (e.g., *blue* printed in red ink). Participants are instructed to ignore the words' meanings and simply to name the color of ink in which each word appears. Across many studies, reaction times (RTs) suggest it is impossible to comply with the instruction to ignore the words' meanings: naming the color of the ink is faster when it is congruent with a word's meaning than when it is incongruent. When we see the word *blue,* we can't help reading it; when we read it, we can't help activating its core meaning, the concept BLUE. Such data seem to support the claim that concepts—or at least their cores—are activated reflexively (i.e., automatically and maximally; Barsalou 1982), whenever we perceive the words that name them, independent of their processing context (which, in the case of the Stroop task, includes instructions *not* to activate words' meanings).

Yet, a small twist on the classic Stroop task shows this conclusion to be a fallacy. Besner and colleagues (1997) posited that the presence of congruent Stroop trials (which typically compose up to 50 percent of the trials in an experiment) were largely responsible for what appeared to be an irrepressible color-meaning interference effect. According to Besner and colleagues, seeing that the ink color matches the meaning of the word on many of the trials reinforces participants' expectation that these elements of the stimulus should match, creating or enhancing the RT disadvantage found for the trials on which they mismatch. If so, then the Stroop effect is, itself, a context effect: congruent trials establish a processing context that makes the incongruent trials especially difficult. To test this account of Stroop interference, rather than comparing RTs for incongruent trials to congruent trials, Besner and colleagues (1997) compared incongruent trials to neutral trials. In one condition, participants named the ink colors for color words printed in the wrong color (e.g., *blue* printed in red ink), and in the other (color-neutral) condition they named the ink colors for meaningless nonwords (e.g., *blat* printed in red ink). Besner and colleagues found the usual large RT disadvantage when responses to incongruent color words were compared to congruent color words (103 millisecond difference), but this effect was reduced to one-third of its size when incongruent color words were compared to color-neutral nonwords (34 millisecond difference). In another condition, Besner and colleagues found the Stroop effect

was eliminated completely when only one of the letters in each word or nonword was colored, rather than all the letters.

In principle, it might be possible to explain away any results showing a failure to activate putatively "core" aspects of a word's meaning by positing that those aspects must not be part of the true core. But this strategy is fraught: How can we ever know what information is in a putative "core" (see section 19.4)? Furthermore, it is hard to imagine any theory of conceptual cores that does not hold "blueness" to be at the core of the concept BLUE. As such, it seems hard to imagine any theory of conceptual cores or default word meanings that can accommodate evidence that blueness is not reflexively activated when people read the word *blue*. Meaning is not a reflex.

19.4   Dissolving the Distinction between *Core* and *Periphery*

We argue that there are no cores to CC&Ms (cf. Margolis and Laurence 2008). The more frequently a piece of information is activated in response to a cue, and the more narrow the range of likely contexts, the more core-like that information may appear. Setting aside experimental results like those of Besner and colleagues (1997) reviewed above, a priori there is no principled way to distinguish an ad hoc concept's *core* from its *periphery.*

As Wittgenstein (1953) illustrated, if we really look for cores, we realize there is nothing to find. It is impossible to identify necessary or sufficient features that could serve as the invariant core of a common artifact concept like GAME, or even of a more constrained-seeming concept like NUMBER (Wittgenstein 1953, sections 66–100). Given a moment's reflection, it becomes clear that even those properties that seem the most likely candidates for core-hood are context dependent. What are the core features, for example, of TIGER? Surely *living* (but what if it's dead?); OK, *animate* (but what if it's a toy tiger?); *striped* (but what if it's bleached, or painted, or selectively bred?); *fierce* (but what if it's tame or lame?); *large* (but what if it's a baby, a dwarf, or a scale model?) If we cannot locate the invariant core of simple concrete ideas like TIGER, how can we hope to find the core of more abstract ideas like LOVE and JUSTICE that are notoriously vague and context dependent? To paraphrase Wittgenstein, there is no way to draw a boundary around a region of property space that includes all and only "core" properties and is valid for all instantiations of a concept or word meaning.

Some concepts like TRIANGLE, ODD NUMBER, and GRANDMOTHER may appear to be definable in terms of certain core, necessary, and sufficient conditions (Armstrong, Gleitman, and Gleitman 1983). Yet, even when people know the conditions and can articulate them, their classification behavior is not predicted by these conditions. For example, although people can articulate the definitions of odd and even numbers, their categorization behavior—both timed and untimed—shows that a number like 798 is *more odd* than a number like 400. Participants reliably miscategorize nonprototypical even numbers like 798 as odd numbers, misclassify scalene triangles as nontriangles, and

believe that grandmothers with fewer grandchildren have a lower chance of winning a contest whose only criterion is being a grandmother (Lupyan 2013). Thus, even though people's explicit definitions of notions like *even number* may be core-like and invariant (e.g., an even number is one that yields an integer when divided by two), their actual categorization behavior cannot be explained by concepts with a fixed core. The apparently crisp definitions available to people metacognitively are not reflected in their cognition.[3]

Arguably, if one averages all the various instantiations of TIGER or GRANDMOTHER over time, people, and contexts, a conceptual core could be said to emerge at their intersection. But this intersection (which could be thought of as a point or as a probability density function in some state space) is a figment of the cognitive scientist's imagination. There is an old joke about three statisticians who go duck hunting. When a duck flies across the horizon, the first statistician shoots one meter above the duck, and the second shoots one meter below the duck. The third statistician calls out, "We got him!" Universal, invariant, conceptual cores can be said to "exist," created by averaging over everyone's idiosyncratic neurocognitive representations, in the same sense that those statisticians can be said to have bagged that duck.

In contrast with theorists who posit that conceptual cores are always activated (e.g., Barsalou 1982; Machery 2009), some theorists posit a kind of core that serves as a default representation that gets activated "most of the time" (Prinz 2002, 157), "when no context is specified" (Prinz 2002, 154). We argue below that a context is *always* specified (though not always explicitly), that the specifics of the context are always in flux (especially when the internal neurocognitive context is taken into account), and that myriad aspects of the internal and external context shape the representations that people form, on multiple time scales.

19.5   Variability of CC&Ms on Three Time Scales

The proposal that all CC&Ms are constructed ad hoc is motivated by the observation that CC&Ms are inherently variable: much more variable than they appear on most characterizations in cognitive science—a field where "representations" were once implemented in punch cards. Unlike the invariant representations in a punch card

3. Explaining these effects as being products of a peripheral "identification procedure" (Armstrong, Gleitman, and Gleitman 1983) hits a dead end when all aspects of categorization of formal concepts are shown to be subject to graded typicality effects (Lupyan 2013). As articulated by Greg Murphy, "At the end of the day, the core serves only as a security blanket to make linguists and philosophers feel better about their concepts, but the behavior all seems to be controlled by the identification procedure. So, if you got rid of the core, your theory's explanatory power would be equal to that of one with a core" (personal communication with G. L., August 13, 2013).

stack, representations in the human mind vary on at least three partly overlapping time scales: (a) from one microsecond to the next within a given instantiation (what we will call *activation dynamics*), (b) from one instantiation to the next within an individual as a function of the *local context*, and (c) from person to person and group to group as a function of people's experiential history (*experiential relativity*; Casasanto and Henetz 2012).

19.5.1   Activation Dynamics

Thinking happens in brains, and brains are always changing. As a consequence, thoughts are always changing. When someone reads the word *dog*, sees a picture of a dog, or sees a dog, what happens after the eyes send the retinal image to V1? When can it be said that they have *seen* the dog or activated the concept DOG, and where in the brain do these acts of perception and conception occur? As information flows from V1, it disperses: some follows the dorsal *where* pathway into the parietal cortex, carried primarily by fast magnocellular connections. Other information follows the ventral *what* pathway carried by the relatively slow parvocellular system.

Even though the new visual information has passed through nearly all the responsive areas within the first 150 milliseconds of this initial "forward sweep" through the brain, visual cortex neurons keep firing at increased rates for hundreds of milliseconds more. During this continued firing, lower-level neurons are responding to feedback from higher brain areas, presumably coordinating the activity in visual areas with distinct receptive fields, composing a coherent percept through dynamic feed-forward, feedback, and lateral connections. The visual brain reconfigures itself continuously throughout the act of seeing (e.g., Lamme and Roelfsema 2000). In parallel, the evolving activity in the visual system engages frontally mediated retrieval processes that incrementally activate relevant information in posterior areas involved in episodic and semantic memory processes.

Where and when, then, is a dog percept or a dog concept activated? The act of perceiving and conceptualizing a dog is not circumscribed, spatially or temporally. The distribution of neurocognitive representations over both time and space makes it impossible to delineate either a moment or an anatomical location in which CC&Ms reside. The perceiver's neurocognitive response to the *dog* stimulus is continuous with their prior neurocognitive activity and with their responses to the context in which the stimulus occurred; it fades gradually into their response to subsequent endogenous and exogenous cues (Spivey 2007).

19.5.2   Local Context

In cognitive psychology, researchers generally try to eliminate effects of context as much as possible, in order to isolate their processes of interest. It is easy to imagine, for example, that by presenting participants with single words or sentences, in randomized

El

order, it should be possible to observe the concepts they "correspond to" in their pure, default form. But there are at least two major problems with this assumption.

The first problem is that, outside the lab, people don't just rely on the cues they're receiving in the experiments to construct an understanding of them; they rely on these cues-in-context. Clark (1997) illustrates how the representations people construct in response to linguistic cues depend critically on their physical and social environments. Even a simple utterance like "I'm hot" may be impossible to interpret in any way approximating its intended meaning unless the listener knows who is saying it and where, when, why, and to whom it is being said. Depending on these factors, this utterance could mean, for example, that the speaker is (1) physically warm, (2) sexually aroused, (3) stolen, (4) radioactive, or (5) on a lucky streak. It is likely that, in the lab, participants might reliably converge on the literal "physical warmth" interpretation of this utterance—but this could be a response to the pragmatics of the experimental context, and could mask other ordinary processes of meaning construction.

There is a second problem with the assumption that we can study CC&Ms in a "neutral context": Even in the most sterile of experimental settings, participants are immersed in a physical context (that of the university psych lab), a social context (interacting with the experimenter, playing the role of the subject, etc.), and their internal biological and neurocognitive context, all of which have consequences for the representations they form. These implicit aspects of the context mold people's minds, in real time, all the time.

A growing catalog of studies at the intersection of metaphor research, emotion, and judgment and decision making illustrates how the incidental physical context affects thinking and language understanding. For example, the spatial metaphor "to move *forward* in time" is ambiguous: it can mean to move earlier or later, depending on the comprehender's spatial perspective. In one study, people's understanding of *forward* was found to vary as a function of their spatial location in a cafeteria line: participants near the end of the line interpreted *forward* to mean earlier in time, whereas those at the front of the line (who had just moved forward through space) thought the same word meant later (Boroditsky and Ramscar 2002). In other studies, the physical context affects memory retrieval. Participants retrieve positive memories more efficiently when they are sitting up tall, and negative memories more efficiently when they are slumping down (Riskind 1983). Likewise, given the same set of memory prompts like "tell me about something that happened last summer," people are more likely to retrieve happy memories when they are assigned to move objects upward during retrieval, and more likely to retrieve sad memories when assigned to move objects downward (Casasanto and Dijkstra 2010). Our actions form one dynamic aspect of the context that shapes our thoughts.

El

Quirks of the physical context affect aspects of the self that people normally think of as stable, such as their morality or their political orientation. In one study, participants in a room with slightly dim lighting were more likely to cheat than those in a room with brighter lighting (Zhong, Bohns, and Gino 2010). This was true even though their cheating was not actually less detectable. In another study, participants were seated in an office chair with one wheel removed, forcing them to lean to the left or the right, while completing a political attitudes questionnaire. Participants who leaned to the left (literally) were more sympathetic to liberal views than those who were induced to lean to the right (Oppenheimer and Trail 2010).

Context is ubiquitous and varied, and so are its effects on the neurocognitive representations people form. Most of the time, the effects of local context on our mental functions are, paradoxically, both dramatic and nearly invisible to us—like the effect of the air around us on bodily functions.

19.5.3   Experiential Relativity

Some aspects of context are more enduring than those reviewed above, such as the language spoken in people's community, their cultural conventions, and the peculiarities of their bodies. Patterns of linguistic, cultural, and bodily experience influence people's minds; people with different patterns of experience tend to think differently, in predictable ways. Together, the influences of these separable streams of experience can be called *experiential relativity* effects (Casasanto and Henetz 2012).

19.5.3.1   Linguistic Relativity   *Linguistic relativity*, the idea that people who speak different languages think differently as a consequence, was once rejected for lack of empirical support. But accumulating evidence suggests that lexical and grammatical differences across languages cause their speakers to form systematically different representations, in fundamental cognitive domains such as *time* (e.g., Boroditsky et al. 2001), *space* (Majid et al. 2004), *motion* (Papafragou, Hulbert, and Trueswell 2008), and *color* (Regier and Kay 2009). Beyond influencing high-level language-mediated thinking, experience using language can also influence low-level perceptual and motor representations, as evidenced by nonlinguistic psychophysical experiments (Casasanto 2008; Dolscheid et al. 2013; see also Lupyan and Spivey 2010a, 2010b; Lupyan and Ward 2013; Thierry et al. 2009). Once we give up the idea of fixed concepts that languages map onto, language can be more productively viewed as a system for activating and manipulating CC&Ms (see Lupyan 2012; Willems and Casasanto 2011; Zwaan 2004).

19.5.3.2   Cultural Relativity   Experiential relativity extends beyond language: Patterns of cultural experience and bodily experience also shape the brain and mind, via some

El

of the same mechanisms by which language does. In one example of *cultural relativity*, across cultures people implicitly conceptualize time as flowing either from left to right or from right to left. This difference cannot be attributed to patterns in language (i.e., no language uses "earlier is *left*" metaphors), or physical experience (i.e., it is not the case that earlier events happen on our left and later events on our right in the natural world). Rather, this habit of thinking appears to be established as people use the particular reading and writing system in their culture (Casasanto and Bottini 2014) and interact with culture-specific artifacts like calendars and graphs (Tversky, Kugelmass, and Winter 1991).

Much broader forms of cultural relativity have also been documented. According to Nisbett and colleagues (2001), members of East Asian and Western cultures have predictably different cognitive styles: whereas Westerners (Europeans and Americans) tend to think individualistically, Easterners (Chinese, Japanese, and Koreans) think more collectivistically, in terms of relationships among members of a society and between people and their environment. Early tests of this proposal seemed too poetic to convince some scientists. For instance, when asked to describe an underwater scene, American participants were likely to start off by mentioning the most prominent individual fish ("there's a big fish …"). By contrast, Japanese participants began by describing the surroundings ("there's a pond …"). Yet, on a skeptical interpretation, such results could merely show that Americans and Japanese people describe things differently, not that they perceive or conceptualize them differently.

Further studies challenged this skeptical position. In one study, Japanese and Americans were shown a box with a vertical line inside of it. They were then shown a second box of a different size and asked to draw a vertical line inside it that matched the one in the first box. Half of the time, participants were told to make the line "the same" as the original, meaning the same absolute length (absolute condition). The other half of the time, they were told to draw a line that was the "same" length as the first in proportion to the surrounding box (relative condition). Results showed that Americans were more accurate in the absolute task, which required focusing on an individual object and ignoring its surroundings, but Japanese participants performed better on the relative task, which required perceiving and remembering an object in its context (Kitayama et al. 2003).

The conclusions about the extent of cross-cultural differences in thinking that are supported by these laboratory tests, for which researchers prioritized experimental control and sought to test cognitive mechanisms, appear dwarfed by field anthropologists' reports of differences in culture-specific practices and attitudes. For example, Henrich and colleagues (2010) report the following difference between cultures:

In the tropical forests of New Guinea, the Etoro believe that for a boy to achieve manhood he must ingest the semen of his elders. This is accomplished through ritualized rites of passage that

require young male initiates to fellate a senior member. … In contrast, the nearby Kaluli maintain that male initiation is only properly done by ritually delivering the semen through the initiate's anus, not his mouth. The Etoro revile these Kaluli practices, finding them disgusting. (61)

To many readers, these rituals may seem immoral; to the Etoro and the Kaluli, apparently they do not, suggesting that concepts like morality and manhood are not universal; they are shaped by specifics of our cultures.

19.5.3.3 Bodily Relativity Beyond linguistic and cultural relativity, different bodily experiences produce *bodily relativity* effects (Casasanto 2011, 2014). For example, right- and left-handers, who typically perform actions like throwing and writing with different hands, also use premotor regions in opposite hemispheres to represent the meanings of the verbs that name these actions (Willems, Hagoort, and Casasanto 2010). Contra all previous models of language in the brain (including those that posit a role for the motor system in constructing verb meanings,;e.g., Pulvermüller 2005), right-hemisphere motor areas are activated by action verbs in people who use these areas to perform common manual actions, demonstrating how small differences in habitual motor actions can induce gross changes in cortical representations of language.

Bodily differences can also affect people's conceptualizations of things they can never see or touch, including their notions of good and bad. Implicitly, right-handers associate positive qualities like honesty, kindness, and intelligence with the right side of space, but left-handers associate positive things with the left, in spite of patterns in language and culture that link *good* with *right.* When asked to decide which of two products to buy, which of two job applicants to hire, or which of two alien creatures looks more trustworthy, right- and left-handers tend to respond differently: right-handers tend to prefer the product, person, or creature presented on their right side, but left-handers prefer the one on their left (Casasanto 2009). This pattern persists even when people make judgments orally, without using their hands to respond. Children as young as five years old already make evaluations according to handedness and spatial location, judging animals shown on their dominant side to be nicer and smarter than animals on their nondominant side (Casasanto and Henetz 2012).

Beyond the laboratory, the association of *good* with the dominant side can be seen in left- and right-handers' spontaneous speech and gestures. In the final debates of the 2004 and 2008 U.S. presidential elections, positive speech was more strongly associated with right-hand gestures and negative speech with left-hand gestures in the two right-handed candidates (Bush, Kerry), but the opposite association was found in the two left-handed candidates (McCain, Obama; Casasanto and Jasmin 2010). In a simulated election, left-handers were more likely than right-handers to vote for a candidate

El

listed on their "good side" of a ballot, suggesting these body-based implicit associations could have significant influences on real-world decisions (Kim, Krosnick, and Casasanto 2014).

19.5.3.4 Flexibility of Experiential Relativity Effects Importantly, the effects of experiential relativity are enduring but also flexible. Habits of mental representation can be changed rapidly when people are exposed to new patterns of linguistic, cultural, or bodily experience. For example, English speakers, who normally conceptualize time as a horizontal line, can be induced to think about time vertically after about twenty minutes of exposure to Mandarin-like up-down metaphors (Boroditsky 2001), or to think about time in terms of three-dimensional space after exposure to Greek-like metaphors that express duration in terms of volume rather than length (Casasanto 2008). Exposing members of a left-to-right-reading culture to a few minutes of mirror-reversed writing can *reverse* the flow of time in their minds, as indicated by RT tests of implicit space-time associations (Casasanto and Bottini 2014). Handedness-based judgments about the abstract ideas of goodness and badness can also be completely reversed, after long-term changes in functional handedness due to unilateral stroke, and even after short-term motor training in the laboratory (Casasanto and Chrysikou 2011).

Are these demonstrations of rapid retraining a problem for interpreting longer-term experiential relativity effects? Only if one adheres to a traditional view of concepts and assumes that "true" relativity effects should result in alterations to concepts' stable cores. If all concepts are constructed ad hoc, in response to cues in context, both long-term and transient relativity effects are demystified. Language, culture, and the body are ubiquitous parts of the context in which people use their minds. To the extent that these aspects of context are constant over time, the mental representations people form will tend to appear constant, exhibiting commonalities across instantiations. To the extent that these aspects of the context vary, the representations people form may vary accordingly (Casasanto 2011, 2014).

19.6 The Illusion of Invariant CC&Ms

The patterns of neurocognitive activity that cognitive scientists characterize as concepts, categories, and word meanings vary according to the contexts in which people instantiate them: from millisecond to millisecond, instantiation to instantiation, individual to individual, and group to group. Why, then, in the face of this pervasive and demonstrable variability, do people have the illusion that CC&Ms are generally stable, across time and across people? We briefly sketch three sources of this illusion.

19.6.1   Conceptual Change Blindness

Observers are often unaware of changes in the visible world around them from one moment to the next, giving rise to phenomena like change blindness. For example, when a participant's face-to-face conversation with an experimenter is briefly interrupted, they often fail to notice changes that occur—even enormous changes—like having their interlocutor vanish and be replaced by a different person (Simons and Levin 1998).

The illusion of stable CC&Ms may arise, in part, from what we will call *conceptual change blindness*. Most of the time, people probably do not introspect on the contents of their ordinary mental representations. But even if they tried, they might be unable to notice, for example, the many ways in which representations cued by a given word differ from one instantiation to the next.

This suggestion is consistent with evidence that, although people generally think they understand how common devices work (e.g., a lock, a zipper), their actual understanding is surprisingly shallow (Keil 2003). As Keil argues, our skeletal knowledge appears more fleshed out than it is, in part because people enrich their theories of how things work on the fly, incorporating information that the current context makes available, ad hoc.

The notorious fallibility of eyewitness memory, in particular people's willingness to substitute incorrect details of a witnessed event for correct ones (e.g., Loftus 1979), provides another line of support for the existence of conceptual change blindness. If conceptual change blindness is anything like visual change blindness, people may intuit invariance even when they are faced with enormous variability in their own mental representations and others'.

19.6.2   The Power of Words

One factor that may contribute to conceptual change blindness, and more broadly to the illusion of invariant CC&Ms, is the use of words to label our ideas—particularly the use of nouns (Barsalou, Wilson, and Hasenkamp 2010). As William James (1890) observed:

Whenever we have a word … to denote a certain group of phenomena we are prone to suppose a substantive entity existing beyond the phenomena, of which the word shall be the name (Barrett, Mesquita, and Smith 2010, 195).

Words encourage psychological essentialism: If we have a word, it is tempting to assume that there must be some observer-independent category of entities that the word names (Barrett, Mesquita, and Smith 2010; Waxman 2004). Verbal labels may help speakers construct representations that are, indeed, more stable across time and across individuals than representations constructed without labels (Lupyan, Rakison,

and McClelland 2007; Lupyan and Thompson-Schill 2012), but verbal labeling also contributes to the *illusion* that different representations activated by the same label are identical to one another.

19.6.3   The Fallacy of Shared Concepts and Mutual Understanding

If CC&Ms are not stable across instances and individuals, how do people ever understand each other? The short answer is: they don't—at least not as well as they might think they do.

On traditional theories of language processing, the result of the comprehension process is a detailed, accurate representation of the input. In reality, however, the representations people form appear to be shadowy reflections of the input, merely "good enough" to achieve communicative goals, at least partially, at least much of the time (Ferreira et al. 2002). Comprehenders typically overlook semantic anomalies in sentences like "the authorities had to decide where to bury the *survivors*." Likewise, they often miss syntactic complexities of sentences like "the dog was bitten by the man," interpreting this unlikely scenario as sensible.

Comprehension failures in these carefully constructed cases may reveal the kind of good-enough processing that people engage in routinely. Rather than a precise, complete, and fully combinatorial process, language use is a process of constructing ad hoc representations cued by words and constructions in context, which may often be shallower, more error prone, and more context dependent than language researchers have suspected (Clark 1996, 1997).

Moreover, successful communication should not be interpreted as evidence that interlocutors shared matching CC&Ms. Suppose, for example, that a speaker says "I saw Julio yesterday." The speaker's memory of "seeing Julio" may involve seeing him in the morning, whereas the listener activated a representation of seeing Julio in the afternoon. This discrepancy may be inconsequential and go unnoticed: we suspect this is the case for innumerable discrepancies between what speakers say and what listeners understand. Understanding is always incremental, is partial, is usually full of errors that go unnoticed, and never represents an exact match between speaker and listener.

In some cases, a match between the speaker's and listener's neurocognitive representations would be undesirable; communicative success often depends on *complementarity* between interlocutors' mental representations. For example, if a speaker says, "I'm freezing!" this speech act would be unsuccessful if what the listener represents is also "I am freezing," or even "He is freezing." It would be a success if this utterance results in the listener representing the intention, "I think I'll close the window" (Grice, Cole, and Morgan 1975; Van Ackeren et al. 2013). Communication is successful to the extent that communicative goals are achieved despite the inevitable mismatches between interlocutors' neurocognitive representations.

19.7   Conclusions

In the effort to explain how people think and communicate, cognitive science has created some of its own hardest problems: mistaking human inventions for natural kinds, and then struggling to explain their "true" nature. Common assumptions about CC&Ms made by cognitive scientists and laypeople are at odds with what is known about the brain and behavior. By rethinking these constructs, it is possible to dissolve some of the classic quandaries at the heart of cognitive science, and to frame a new set of challenges for research on language and mind.

According to the AHC framework, all concepts, categories, and word meanings are constructed ad hoc and differ from one instantiation to the next, within and between individuals and groups. CC&Ms are inseparable from their contexts, and are shaped by the contexts in which they are instantiated on every time scale, from the millisecond to the lifetime. Even in "fictive preparations" of CC&Ms in the psychology lab, they are not context independent; experiments occur in specific physical and social contexts. More generally, people's memories, languages, cultures, and bodies are ubiquitous parts of the context in which they use their minds.

If future research can demonstrate the kind of invariance in CC&Ms that is commonly assumed, AHC will have to be revised or abandoned. We have argued, however, that past research that has been interpreted as evidence of invariance in CC&Ms is not convincing; despite people's intuitions, invariance is an illusion. If this is true, then a central goal of research on language and cognition should be to elucidate the fleeting, idiosyncratic neurocognitive representations that people *actually use* for thinking and communicating—and to explain how apparent stability emerges from pervasive variability.

References

Allott, N., and M. Textor. 2012. Lexical pragmatic adjustment and the nature of ad hoc concepts. *International Review of Pragmatics* 4 (2): 185–208.

Altmann, G., and Y. Kamide. 1999. Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition* 73 (3): 247–264.

Armstrong, S. L., L. R. Gleitman, and H. Gleitman. 1983. What some concepts might not be. *Cognition* 13 (3): 263–308.

Barrett, L. F., B. Mesquita, and E. R. Smith. 2010. The context principle. In *The Mind in Context*, ed. B. Mesquita, L. F. Barrett, and E. R. Smith, 1–21. New York: Guilford Press.

Barsalou, L. W. 1982. Context-independent and context-dependent information in concepts. *Memory & Cognition* 10:82–93.

El

Barsalou, L. W. 1983. Ad hoc categories. *Memory & Cognition* 11:211–227.

Barsalou, L. W. 1987. The instability of graded structure: Implications for the nature of concepts. In *Concepts and Conceptual Development*, ed. U. Neisser, 101–140. Cambridge: Cambridge University Press.

Barsalou, L. W., C. D. Wilson, and W. Hasenkamp. 2010. On the vices of nominalization and the virtues of contextualizing. In *The Mind in Context*, ed. B. Mesquita, L. Feldman-Barrett, and E. Smith, 334–360. New York: Guilford Press.

Besner, D., J. A. Stolz, and C. Boutilier. 1997. The Stroop effect and the myth of automaticity. *Psychonomic Bulletin & Review* 4 (2): 221–225.

Boroditsky, L. 2001. Does language shape thought? English and Mandarin speakers' conceptions of time. *Cognitive Psychology* 43 (1): 1–22.

Boroditsky, L., and M. Ramscar. 2002. The roles of body and mind in abstract thought. *Psychological Science* 13 (2): 185–188.

Casasanto, D. 2008. Who's afraid of the Big Bad Whorf? Cross-linguistic differences in temporal language and thought. *Language Learning* 58 (1): 63–79.

Casasanto, D. 2009. Embodiment of abstract concepts: Good and bad in right- and left-handers. *Journal of Experimental Psychology: General* 138 (3): 351–367.

Casasanto, D. 2011. Different bodies, different minds: The body specificity of language and thought. *Current Directions in Psychological Science* 20 (6): 378–383.

Casasanto, D. 2014. Bodily relativity. In *Routledge Handbook of Embodied Cognition*, ed. L. Shapiro, 108–117. New York: Routledge.

Casasanto, D., and R. Bottini. 2014. Mirror reading can reverse the flow of time. *Journal of Experimental Psychology: General* 143(2): 473–479.

Casasanto, D., and E. G. Chrysikou. 2011. When left is "right": Motor fluency shapes abstract concepts. *Psychological Science* 22 (4): 419–422.

Casasanto, D., and K. Dijkstra. 2010. Motor action and emotional memory. *Cognition* 115 (1): 179–185.

Casasanto, D., and T. Henetz. 2012. Handedness shapes children's abstract concepts. *Cognitive Science* 36:359–372.

Casasanto, D., and K. Jasmin. 2010. Good and bad in the hands of politicians. *PLoS ONE* 5 (7): e11805.

Chomsky, N. 2002. *On Nature and Language*. Cambridge: Cambridge University Press.

Churchland, P. S. 1986. *Neurophilosophy: Toward a Unified Science of the Mind/Brain*. Cambridge, MA: MIT Press.

Clark, H. H. 1996. *Using Language*. Cambridge: Cambridge University Press.

Clark, H. H. 1997. Dogmas of understanding. *Discourse Processes* 23 (3): 567–598.

DeLong, K. A., T. P. Urbach, and M. Kutas. 2005. Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience* 8 (8): 1117–1121.

Dikker, S., H. Rabagliati, T. A. Farmer, and L. Pylkkänen. 2010. Early occipital sensitivity to syntactic category is based on form typicality. *Psychological Science* 21 (5): 629–634.

Dolscheid, S., S. Shayan, A. Majid, and D. Casasanto. 2013. The thickness of musical pitch: Psychophysical evidence for linguistic relativity. *Psychological Science* 24 (5): 613–621.

Dunham, Y., and M. R. Banaji. 2010. Platonic blindness and the challenge of understanding context. In *The Mind in Context*, ed. L. Feldman-Barrett, B. Mesquita, and E. Smith, 201–213. New York: Guilford.

Elman, J. L. 2004. An alternative view of the mental lexicon. *Trends in Cognitive Sciences* 8 (7): 301–306.

Elman, J. L. 2009. On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science* 33 (4): 547–582.

Evans, V. 2009. *How Words Mean.* Oxford: Oxford University Press.

Ferreira, F., V. Ferraro, and K. Bailey. 2002. Good-enough representations in language comprehension. *Current Directions in Psychological Science* 11:11–15.

Grice, H. P., P. Cole, and J. Morgan. 1975. Logic and conversation. In *Syntax and Semantics,* vol. 3: *Speech Acts,* ed. P. Cole and Jerry L. Morgen, 45–47. New York: Academic Press.

Hampton, J. A. 2012. Thinking intuitively: The rich (and at times illogical) world of concepts. *Current Directions in Psychological Science* 21 (6): 398–402.

Henrich, J., S. J. Heine, and A. Norenzayan. 2010. The weirdest people in the world? *Behavioral and Brain Sciences* 33 (2–3): 61–83, discussion 83–135.

Jackendoff, R. S. 2002. *Foundations of Language: Brain, Meaning, Grammar, and Evolution.* Oxford: Oxford University Press.

James, W. 1890. *Principles of Psychology.* Vol. 1. New York: Holt.

Katz, J., and P. Postal. 1964. *An Integrated Theory of Linguistic Descriptions.* Cambridge, MA: MIT Press.

Keil, F. 2003. Folkscience: Coarse interpretations of a complex reality. *Trends in Cognitive Sciences* 7 (8): 368–373.

Kim, N., J. Krosnick, and D. Casasanto. In press. Moderators of candidate name order effects in elections: An experiment. *Political Psychology.*

Kitayama, S., S. Duffy, T. Kawamura, and J. T. Larsen. 2003. Perceiving an object and its context in different cultures: A cultural look at new look. *Psychological Science* 14:201–206.

Lamme, V., and P. Roelfsema. 2000. The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences* 23 (11): 571–579.

Langacker, R. 2008. *Cognitive Grammar: A Basic Introduction*. New York: Oxford University Press.

Loftus, E. F. 1979. *Eyewitness Testimony*. Cambridge, MA: Harvard University Press.

Lupyan, G. 2012. What do words do? Towards a theory of language-augmented thought. In *The Psychology of Learning and Motivation*, vol. 57, ed. B. H. Ross, 255–297. London: Academic Press.

Lupyan, G. 2013. The difficulties of executing simple algorithms: Why brains make mistakes computers don't. *Cognition* 129 (3): 615–636.

Lupyan, G., D. H. Rakison, and J. L. McClelland. 2007. Language is not just for talking: Labels facilitate learning of novel categories. *Psychological Science* 18 (12): 1077–1082.

Lupyan, G., and M. Spivey. 2010a. Making the invisible visible: Auditory cues facilitate visual object detection. *PLoS ONE* 5 (7): e11452.

Lupyan, G., and M. Spivey. 2010b. Redundant spoken labels facilitate perception of multiple items. *Attention, Perception & Psychophysics* 72 (8): 2236–2253.

Lupyan, G., and S. L. Thompson-Schill. 2012. The evocative power of words: Activation of concepts by verbal and nonverbal means. *Journal of Experimental Psychology: General* 141 (1): 170–186.

Lupyan, G., and E. J. Ward. 2013. Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences* 110(35): 14196–14201.

Machery, E. 2009. *Doing Without Concepts*. New York: Oxford Univ. Press.

Mahon, B. Z., and A. Caramazza. 2008. A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology, Paris* 102 (1): 59–70.

Majid, A., M. Bowerman, S. Kita, D. B. M. Haun, and S. C. Levinson. 2004. Can language restructure cognition? The case for space. *Trends in Cognitive Sciences* 8 (3): 108–114.

Margolis, E., and S. Laurence. 2008. Concepts. In *The Blackwell Guide to Philosophy of Mind,* vol. 10, ed. S. Stich and T. Warfield, 190–213. Oxford: Wiley-Blackwell.

Nisbett, R. E., K. Peng, I. Choi, and A. Norenzayan. 2001. Culture and systems of thought. *Psychological Review* 108:291–310.

Oppenheimer, D. M., and T. E. Trail. 2010. Why leaning to the left makes you lean to the left: Effect of spatial orientation on political attitudes. *Social Cognition* 28 (5): 651–661.

Osherson, D., and E. Smith. 1981. On the adequacy of prototype theory as a theory of concepts. *Cognition* 9:35–58.

Papafragou, A., J. Hulbert, and J. Trueswell. 2008. Does language guide event perception? *Cognition* 108 (1): 155–184.

El

Prinz, J. 2002. *Furnishing the Mind*. Cambridge, MA: MIT Press.

Pulvermüller, F. 2005. Brain mechanisms linking language and action. *Nature Reviews. Neuroscience* 6 (7): 576–582.

Regier, T., and P. Kay. 2009. Language, thought, and color: Whorf was half right. *Trends in Cognitive Sciences* 13:439–446.

Riskind, J. H. 1983. Nonverbal expressions and the accessibility of life experience memories: A congruence hypothesis. *Social Cognition* 2:62–86.

Rogers, T. T., and J. L. McClelland. 2004. *Semantic Cognition: A Parallel Distributed Processing Approach*. Cambridge, MA: MIT Press.

Simons, D. J., and D. T. Levin. 1998. Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review* 5:644–649.

Smith, L. B., and L. K. Samuelson. 1997. Perceiving and remembering: Category stability, variability and development. In *Knowledge, Concepts, and Categories*, ed. K. Lamberts and D. Shanks, 161–196. Cambridge, MA: MIT Press.

Snedeker, J., and L. Gleitman. 2004. Why is it hard to label our concepts? In *Weaving a Lexicon,* ed. D. G. Hall and S. R. Waxman, 257–294. Cambridge, MA: MIT Press.

Spivey, M. 2007. *The Continuity of Mind*. New York: Oxford University Press.

Stroop, J. R. 1935. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology* 18 (6): 643–662.

Taylor, L. J., and R. A. Zwaan. 2009. Action in cognition: The case of language. *Language and Cognition* 1:45–58.

Thierry, G., P. Athanasopoulos, A. Wiggett, B. Dering, and J.-R. Kuipers. 2009. Unconscious effects of language-specific terminology on preattentive color perception. *Proceedings of the National Academy of Sciences of the United States of America* 106 (11): 4567–4570.

Tversky, B., S. Kugelmass, and A. Winter. 1991. Cross-cultural and developmental trends in graphic productions. *Cognitive Psychology* 23:515–557.

Van Ackeren, M., D. Casasanto, H. Bekkering, P. Hagoort, and S. A. Rueschemeyer. 2012. Pragmatics in action: Indirect requests engage theory of mind areas and the cortical motor network. *Journal of Cognitive Neuroscience* 11 (4): 2237–2247.

Waxman, S. R. 2004. Everything had a name, and each name gave birth to a new thought: Links between early word-learning and conceptual organization. In *Weaving a Lexicon*, ed. D. G. Hall and S. R. Waxman, 295–335. Cambridge, MA: MIT Press.

Weiskopf, D. 2009. The plurality of concepts. *Synthese* 169:145–173.

Willems, R. M., and D. Casasanto. 2011. Flexibility in embodied language understanding. *Frontiers in Psychology* 2 (116): 1–11.

El

Willems, R. M., P. Hagoort, and D. Casasanto. 2010. Body-specific representations of action verbs: Neural evidence from right- and left-handers. *Psychological Science* 21 (1): 67–74.

Wittgenstein, L. 1953. *Philosophical Investigations*. Ed. G. E. M. Anscombe and R. Rhees. Trans. G. E. M. Anscombe. Oxford: Blackwell.

Zhong, C., V. K. Bohns, and F. Gino. 2010. Good lamps are the best police: Darkness increases dishonesty and self-interested behavior. *Psychological Science* 21:311–314.

Zwaan, R. A. 2004. The immersed experiencer: Toward an embodied theory of language comprehension. In *The Psychology of Learning and Motivation*, vol. 44, ed. B. H. Ross, 35–62. New York: Academic Press.

El