Language in Mind

An Introduction to Psycholinguistics

JULIE SEDIVY *University of Calgary*

SINAUER ASSOCIATES, INC. • Publishers • Sunderland, Massachusetts • USA

mind. Recurring patterns across languages might simply reflect events of history. Many different languages can evolve out of a common ancestral tongue, and it wouldn't be surprising to find striking similarities among them merely due to their shared origins. Or, similarities could arise because speakers of two different languages came into regular contact with each other and influenced each other (see **Box 12.1** on the previous page). For instance, modern English offers a captivating record of the many imprints that other languages have left upon it. Such imprints are most obvious in the form of "borrowed" words like *lingerie* (French), *maximum* (Latin), *democracy* (Greek), *marijuana* (Spanish), *kindergarten* (German), or *futon* (Japanese). According to some estimates, as much as 75% of the vocabulary of English originated in other languages, largely Latin and French (Thomason, 2001). But language contact doesn't just affect word borrowings; more structural elements of phonology or syntax can also creep into languages from other tongues.

These historical accounts can be fascinating in their own right. But they don't necessarily tell us that much about how language works in the human mind, or where the boundaries of human language lie. To get a sense of the deeper constraints underlying language, it's important to survey a very large number of languages, placing special emphasis on those patterns that come up again and again even for languages that aren't historically related to each other, and whose speakers have not had contact with each other.

In one ambitious line of work, linguist Matthew Dryer (1992) tested Greenberg's universals against a much larger set of 625 languages, and found that, while some of Greenberg's proposed universals failed to hold up in the larger sample, others did quite well. There seemed to be particularly strong evidence for certain word order correlations. For example, if a language places the verb before the direct object in a sentence, you can bet on the fact that it also has prepositions, which occur *before* their associated noun phrases (also called prepositional objects). English is one such language (let's call it a "Type A" language):

Dimitri	swept	the porch	with	a broom.
	verb	object	preposition	prepositional object

On the other hand, if a language places the verb *after* its object, as does Japanese, the chances are very good that it makes use of *post*positions, which occur *after* the associated noun phrase (let's call this a "Type B" language):

Taroo-	*waboo-*	*deinu*	*-obutta*
Taroo	stick-with	dog	hit

Strong correlations like these have been taken as evidence of a deep cognitive bias. Researchers who argue for the existence of an innate universal grammar have suggested that we have inborn "settings" that constrain the possible word orders that we learn as children—we innately "know" that a language is either Type A or Type B. Other researchers have argued that the strong word order correlations reflect a general processing preference to keep word order rules consistent across different kinds of phrases.

More recently, Michael Dunn and his colleagues (2011) have argued that the word order correlations that characterize the Type A versus Type B languages don't reveal anything especially deep about the nature of human languages; instead, they simply reflect the historical development of languages. To support this argument, Dunn and his colleagues built a statistical model to take into account the lineage histories of languages. If word order correlations are indeed "true" universals that constrain the preferred word orders of human language, we'd expect them to appear independent of language lineage

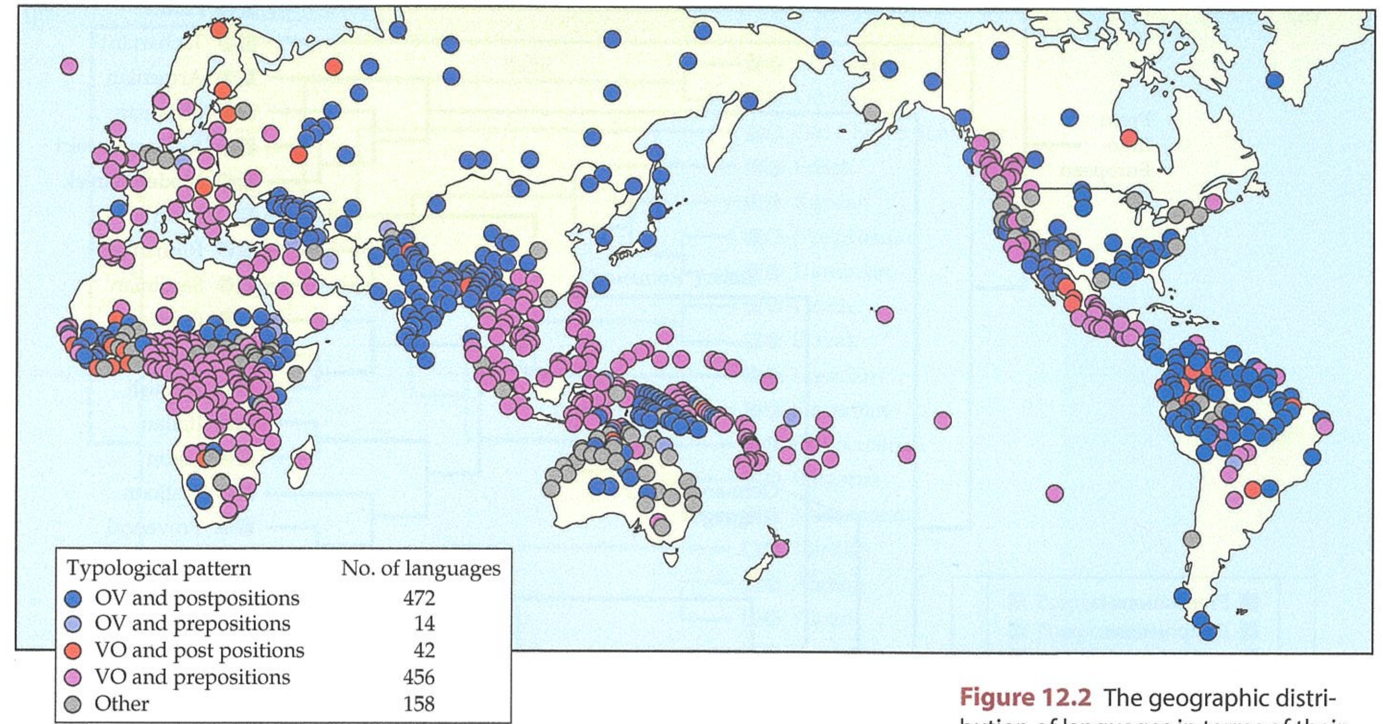

Figure 12.2 The geographic distribution of languages in terms of their ordering of the verb (V) and object (O) and presence of pre- and postpositions. Note how languages cluster geographically by type, suggesting that typological patterns may be due to linguistic lineage or language contact. (From Dryer, 2013c; see http://wals.info/feature/95A#2/14.9/152.8.)

or geography. But Type A and Type B languages divide up rather neatly by geography, as shown in **Figure 12.2**. Furthermore, Type A and Type B word order patterns tend to cluster in languages that are very closely related to each other, as shown in **Figure 12.3**.

The fact that word order type can be largely predicted by a language's historical roots and/or its geographic context makes it hard to rule out the possibility that either a common ancestry or contact between languages is responsible for the close connection between verb-object order and whether a language has prepositions or postpositions. Under this account, there's nothing "deep" about the fact that languages that have the verb-object order also tend to have prepositions rather than postpositions; the correlation simply reflects the fact that languages with a shared history are likely to be similar in a *number* of different ways. If you looked closely enough, you might find many other correlations, none of them particularly deep or meaningful.

Here's an analogy: Suppose you find a correlation between a language's use of tone to distinguish lexical meanings and the likelihood that its speakers use chopsticks as eating utensils. Does this point to some intrinsic connection between lexical tone and chopstick use? That seems rather unlikely. A more plausible interpretation is simply that the speakers of various tone languages share a common cultural and linguistic heritage, which included, among other things, the wielding of chopsticks during dinner.

Are there meaningful constraints on the shape of human languages, constraints that truly reflect something about the human mind? Looking at recurring crosslinguistic patterns and universals can provide some hints (and nowadays researchers can begin to explore interesting patterns with the help of research tools such as the online database of languages found in the World

WEB ACTIVITY 12.1

Variation across languages In this activity, you'll explore the online World Atlas of Language Structures (WALS). This resource identifies many ways in which languages vary from one another, and allows you to see the geographic distributions of various linguistic features.

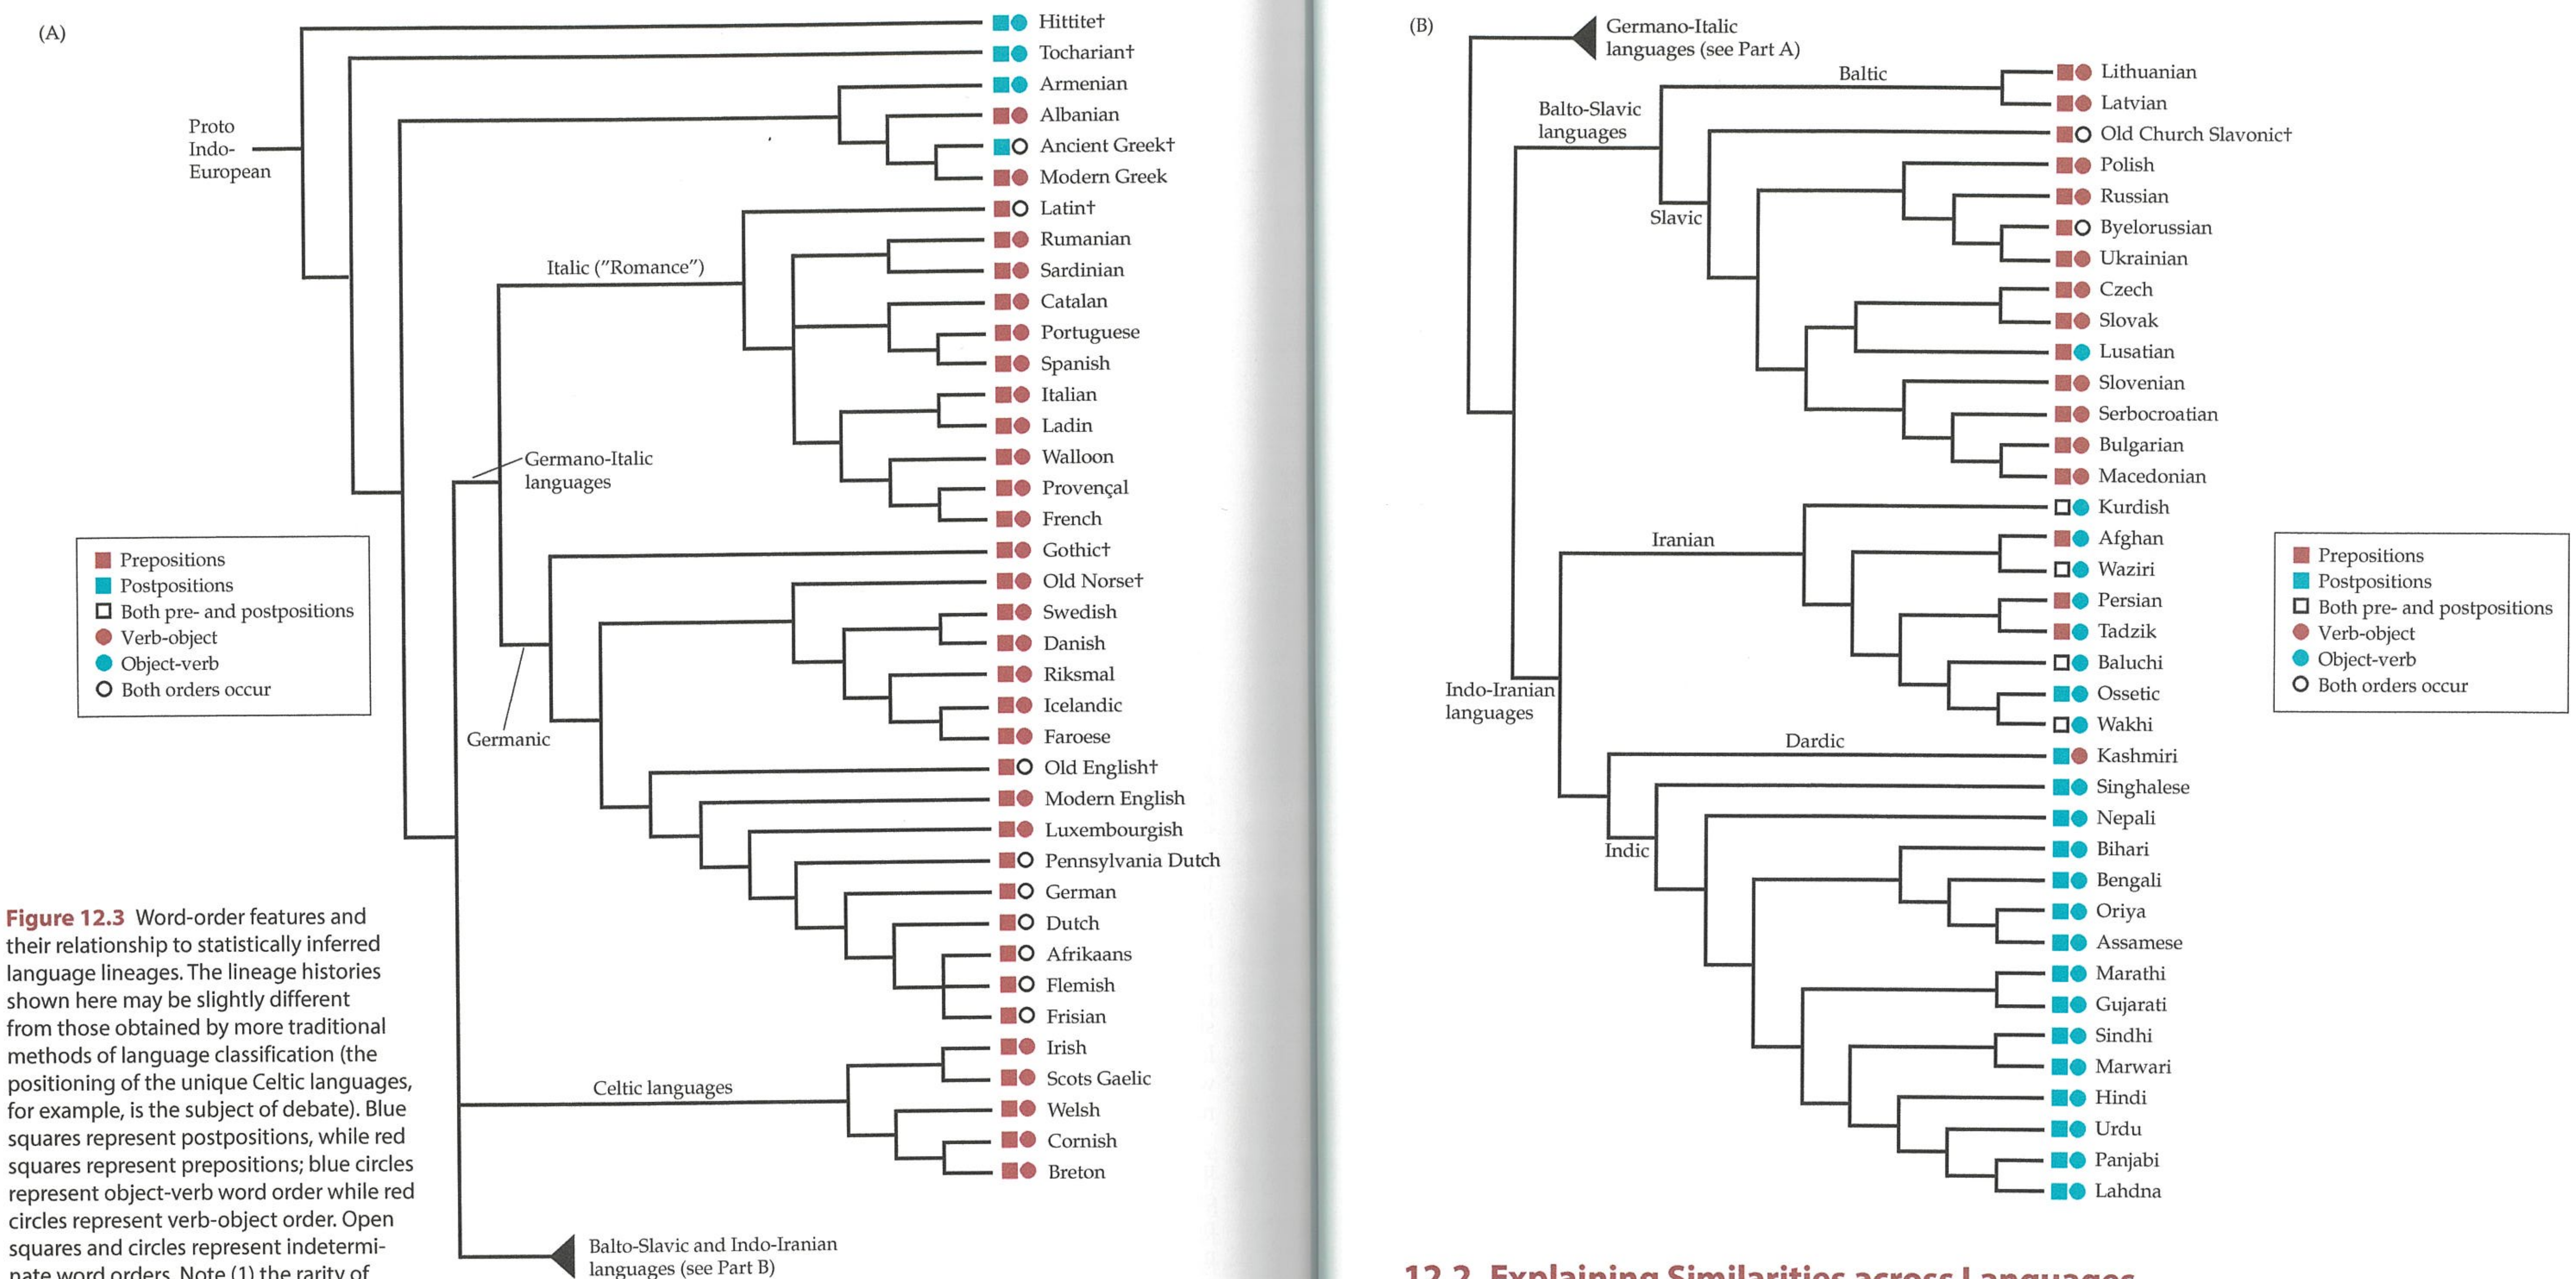

Figure 12.3 Word-order features and their relationship to statistically inferred language lineages. The lineage histories shown here may be slightly different from those obtained by more traditional methods of language classification (the positioning of the unique Celtic languages, for example, is the subject of debate). Blue squares represent postpositions, while red squares represent prepositions; blue circles represent object-verb word order while red circles represent verb-object order. Open squares and circles represent indeterminate word orders. Note (1) the rarity of mixed (red-blue or blue-red) elements; and (2) the degree to which linguistic elements are similar across closely related languages. (Adapted from Dunn et al., 2011.)

Atlas of Language Structures, or WALS). But as we've just seen, there are some remaining challenges in sorting out exactly where these similarities come from and how extensive they are. Luckily, we can work toward this question from another angle: we can start by looking at whether certain patterns are generally easier for people to learn or to produce than others, and by seeing whether there's a connection between the easy patterns and those that are commonly found across languages.

12.2 Explaining Similarities across Languages

Learning biases

Many researchers have argued that the existence of universal or recurring language patterns is evidence that the human mind plays favorites, finding some linguistic forms more "learnable" than others. But as you just saw, it's not always clear whether these crosslinguistic similarities have a cognitive basis as opposed to a historical one. Hence, a number of language scientists are finding more direct ways to search for learning biases.

ticular idea. In that chapter, you read about how speakers' choices can be driven by cognitive pressures; in any particular instance, people will often choose a linguistic form that makes the arduous task of speaking just a little bit easier. Now, what if we zoom out and look at the production of language through the lens of linguistic universals or tendencies? Some provocative questions and predictions quickly arise.

Are some linguistic forms systematically easier to produce than others? If so, we'd expect that, wherever their language permitted, speakers would tend to use the easier forms more often than the more difficult options for expressing the same idea. This simple asymmetry could set in motion a cross-generational language shift: the input to new language learners would be riddled with the easier-to-say forms, while the harder-to-produce structures would be more sparse. We've already seen that new learners of a language have a tendency to over-regularize the input they hear; therefore, the next generation of language users would be prone to exaggerate these statistical differences. Over time, the harder structures might drop out of the language entirely, while the easier structures would be preserved. So we'd predict that if were to look across languages, we'd be more likely to find structures that ease the demands of speaking rather than corresponding structures that put more stress on cognitive resources during speaking.

Let's take a more concrete look at how production pressures could lead to common crosslinguistic patterns. I'll draw on an example from Maryellen MacDonald (2013), one researcher who has argued that production pressures are likely to play an important role in explaining crosslinguistic tendencies. In English, if we want to describe an event that involves two participants—the subject (S) and the object (O) of a verb (V)—we normally use the word order: subject-verb-object (SVO). The SVO order is just one of six possible ways in which these three linguistic units could be combined. But if we look across many of the world's languages, some of the options are wildly more popular than others. The vast majority of languages embrace the solution of placing the subject first; and languages seem to be almost allergic to ordering the object before the subject (see **Figure 12.6**).

Is it possible to tell a story about how production pressures might lead to a bias for placing the subject toward the beginning of a sentence? As discussed in Chapter 9, we know that when a word or phrase is highly accessible, speakers tend to utter it as quickly as possible. The general idea is that as soon as speakers are mentally prepared to utter a word or phrase, they spit it out in order to avoid clogging up working memory while planning the rest of the sentence. We saw that various factors could affect the accessibility of a linguistic unit. For example, shorter phrases were more likely than longer phrases to be uttered early in the sentence. Words could be made more accessible through previous mention, or through priming with a semantically related word. And in some experimental manipulations, characters or objects in a visual scene were made more visually salient through the use of flashing markers—with the result that drawing visual attention to an entity made it more likely that speakers would mention it first.

If, upon planning a sentence, speakers find that it's especially easy to bring to mind the instigators of the action they want to describe, then they should prefer to order subjects first, wherever possible. Some researchers (e.g., Bock et al., 1992) have suggested that animate concepts, such as *dentist* or *woman*, are retrieved from memory more easily than inanimate ones, such as *flower* or *car*. Animacy is very closely linked to subjecthood; notice that whenever an event involves an animate participant and an inanimate entity, the animate participant is almost always the subject—just try coming up with sentences involving the word pairs *dentist/flowers*, *steak/panther*, and *boy/book*.

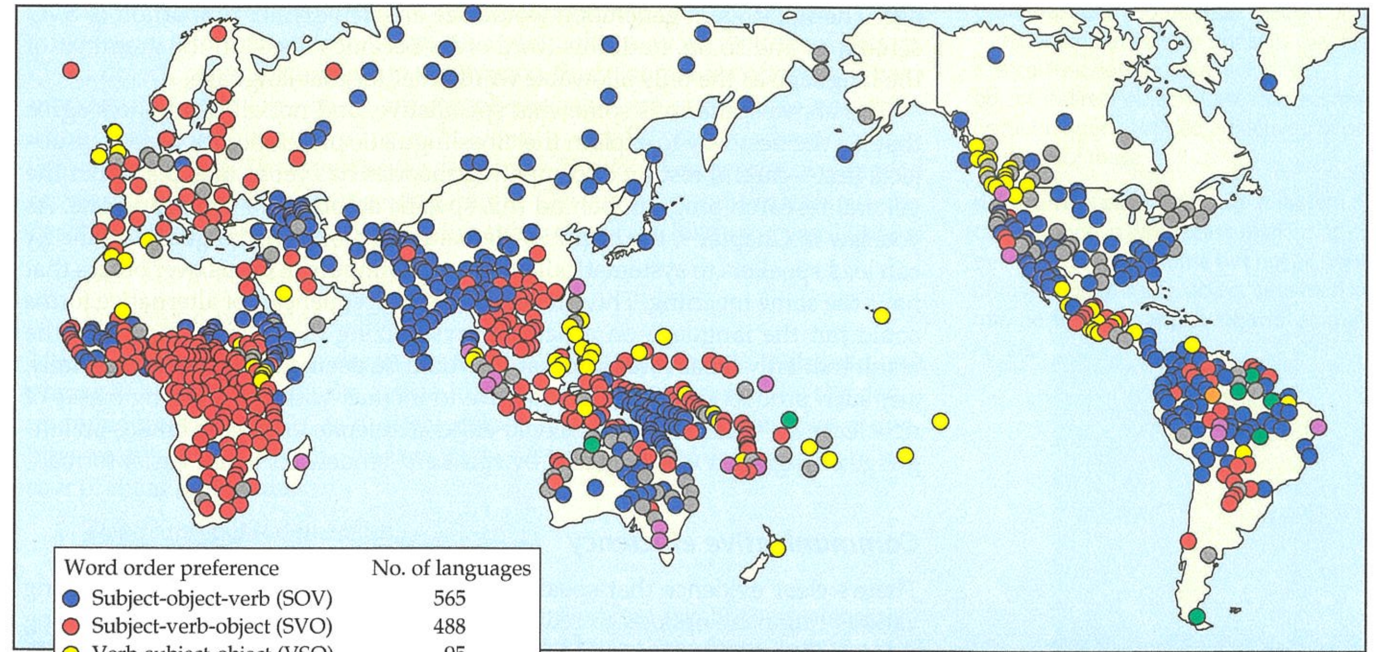

Word order preference	No. of languages
Subject-object-verb (SOV)	565
Subject-verb-object (SVO)	488
Verb-subject-object (VSO)	95
Verb-object-subject (VOS)	25
Object-verb-subject (OVS)	11
Object-subject-verb (OSV)	4
Lacking a dominant word order	189
TOTAL	1377

Figure 12.6 The distribution of all possible orderings of the subject, object, and verb, as drawn from the WALS database. Note the dramatic shortage of languages that place the object before the subject. (Adapted from Dryer, 2013a; see http://wals.info/feature/81A#2/18.0/152.8.)

Indeed, a number of studies have shown that speakers exercise their grammatical options in such a way as to order animate nouns earlier in the sentence than inanimate nouns. In the study by Kay Bock and her colleagues, subjects were more likely to use the passive voice if it resulted in the animate participant being ordered before the inanimate object in the sentence: *The boy was hit by the truck*, rather than *The truck hit the boy*.

Let's imagine that the English language evolved from a previous state in which word order was quite fluid and multiple word orders were allowed, as is the case for a number of existing languages. So, for example, all of the following sentences could have meant the same thing:

The students devoured the free pizza. (SVO)

The students the free pizza devoured. (SOV)

Devoured the students the free pizza. (VSO)

Devoured the free pizza the students. (VOS)

The free pizza the students devoured. (OSV)

The free pizza devoured the students. (OVS)

Because of the pressures on language production just described, the subject-first orders would be easier to produce than others, and hence would be the most likely to be uttered. Let's suppose that the SVO order was the most commonly produced order. (In Web Activity 12.2, you'll have the opportunity to explore why this particular order might be preferred over the other subject-first solutions.) The next generation of learners, who heard the SVO order more often than any of the others, would magnify this bias in their own language us-

It's likely that some conceptual distinctions are universally more salient than others, leading to some predictable ways of structuring the lexicon of any language. For example, in his book *Through the Language Glass*, linguist Guy Deutscher (2010) invites you to imagine coming across an old manuscript that describes a language called "Ziftish." In Ziftish, it turns out, there is a word *bose*, which is used to refer to white roses and all birds except those with red chests. Another word, *rird*, is used for red-chested birds and all roses except white ones. Deutscher wants to know: Do you, the discoverer of this manuscript, take it to be a factual diary of an early explorer? Or a fictional account—perhaps a long-lost sequel to *Gulliver's Travels*? The manuscript reeks of fiction, because it seems deeply implausible that a language would confer words on such unnatural categories.

This example suggests that salient or natural concepts attract words more readily than less natural ones. So what about when languages diverge in how they map concepts onto language? A reasonable hypothesis is that the divergence reflects how important these concepts are for particular communities of speakers.

It's not hard to come up with examples where certain lexical distinctions align neatly with the cultural importance of their corresponding concepts. Surely, a culture where all food is cooked on a spit over the fire or warmed up in the microwave has no need for a specialized vocabulary that distinguishes between words like *sauté, braise, grill, boil, bake, blanch, poach, broil, simmer, fricassee, flambé, steam, fry, caramelize, stew, sear,* and so on. When it comes to such words, necessity is a plausible mother of invention.

Perhaps the most striking example of how culture can drive the invention of words is in the domain of color vocabulary. Even within a single language like English, some speakers feel compelled to specify that a color is *magenta* or *chartreuse*, or to distinguish between *crimson* and *scarlet*, while others are perfectly satisfied with basic color terms like *blue, yellow,* and *red*. But if you think that all languages at least distinguish between the basic colors—much as they'd likely have separate words for general categories like birds or flowers—you'd be wrong. For example, it may surprise you to hear that color terms are virtually absent from the ancient Greek epics of Homer, as noted in 1858 by classics scholar William Gladstone (who is better known for having served four terms as the prime minister of Great Britain). This observation led Gladstone to speculate that the ancient Greeks were color-blind, and that humans have only very recently developed color vision as we have it today. This notion was taken seriously by many scientists of the time, even though some questioned whether the evolution of color vision could have developed in the short span of several thousand years of human history.

But in the twentieth century, as linguists and anthropologists began to comb through the existing languages of the world, they found many languages that didn't have separate words for green versus blue, for example, or even red versus yellow (see Box 12.3). The speakers of these languages weren't color-blind. They could *see* that green and blue were different colors, they just thought it would be odd to call them by different names—just as you might be able to see that two slightly different shades of red are different from each other, but wonder why anyone would need to have different words for them. Speakers of all languages, it turns out, can detect color differences that they don't bother to mark with different words. But language communities can differ widely in the number of distinct color words they feel are necessary. Many of the languages with very small color vocabularies are spoken in non-industrial societies that don't do much manufacturing of objects involving artificial color. In such a

BOX 12.3
Variations in color vocabulary

What could be more basic than color concepts like brown, green, yellow, and red? A survey of color terms across the world's languages shows surprising diversity in the number of color terms that are used in a language. In English, if we consider just those color terms that correspond to single, commonly used words, we have a total of eleven: *black, white, red, yellow, green, blue, purple, brown, orange, pink,* and *gray*. This represents the upper end of the vocabulary size for basic color terms across languages, as shown in Figure 12.7. Many languages make do with as few as three terms to express color.

Some languages use simple color words to make distinctions that we don't make in English. For example, Russian and Greek have separate words for dark and light shades of blue. But many other languages use a single term to name colors that English refers to by different names. For example:

- Yupik (as spoken in Siberia) and Pirahã (Brazil; see Box 6.3) use a single word for green and blue.
- Lele (Chad) and Javaé (Brazil) use a single word for yellow, green, and blue.
- Gunu (Cameroon) and Tacana (Bolivia) refer to red and yellow with a single word.

Some researchers (most notably Brent Berlin and Paul Kay, 1969) have argued that there are certain important color-naming universals reflecting underlying perceptual constraints. As with linguistic structures, some color terms appear to be more common than others. For instance, if a language has only three color terms, the terms tend to separate into these categories: (1) white; (2) red and yellow; and (3) black, green, and blue. And, while it's common for languages to use a single word for green and blue, basic word distinctions between dark and light blue are rare.

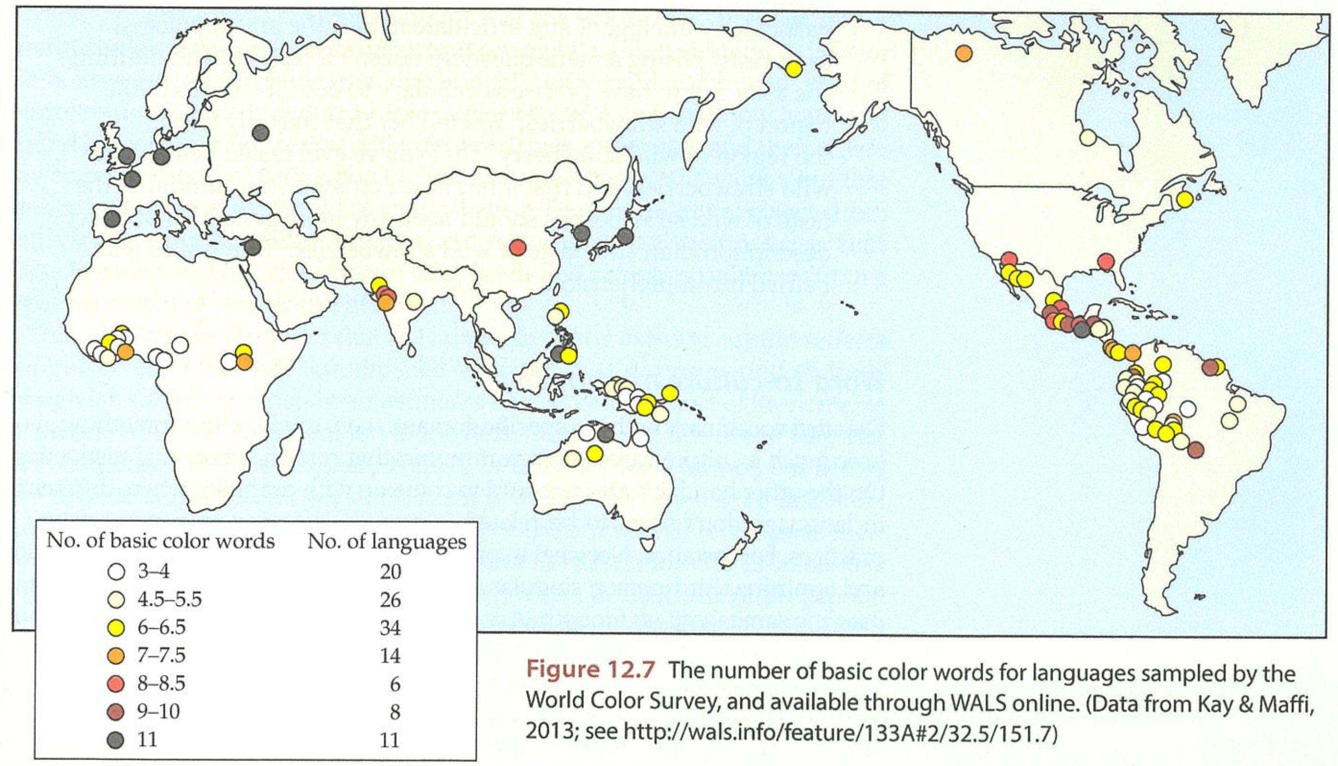

Figure 12.7 The number of basic color words for languages sampled by the World Color Survey, and available through WALS online. (Data from Kay & Maffi, 2013; see http://wals.info/feature/133A#2/32.5/151.7)

society, the color of an object is largely predictable from its inherent nature, so why would you bother to specify it?

To help you imagine what it might be like to live in a culture where a detailed color vocabulary seems unimportant, Guy Deutscher invites you to imagine a

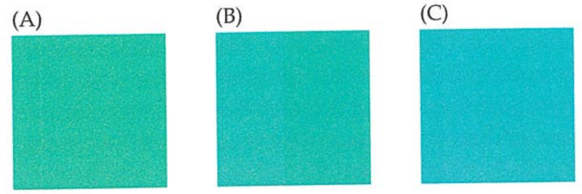

Figure 12.9 Example stimuli used by Kay and Kempton (1984) in the "odd one out" task for color perception. Subjects were asked to identify which of the color samples was more distant from the other two samples than those two were from each other. Speakers of English, but not speakers of Tarahumara, tended to exaggerate the distance between stimuli B and C, which span the English blue/green boundary.

dimensions of hue, brightness, and saturation.) For each of their experimental trials, Kay and Kempton chose three closely neighboring color chips from the Munsell chart. Subjects were told that two of the colors were very similar to each other, but that one was more distant from the other two; the subjects' task was to identify the "odd one out" (**Figure 12.9**). Based on the subjects' responses over a number of trials, the researchers calculated a score of perceived distance between critical pairs of color chips, and compared these to measures of objective distance. They found that the language spoken by subjects did have an effect on the perceived distances between colors. English speakers judged the distance between a "blue" and a "green" chip to be bigger than the distance between two shades of "blue," even if the objective distances were identical. (The boundaries between the English green and blue categories were previously established on the basis of responses from English speakers who did not participate in the "odd one out" task.) On the other hand, Tarahumara speakers did not exaggerate the distance between colors across the green/blue boundary.

Is this convincing evidence that having different labels for green and blue has warped the perceptual system of English speakers? Perhaps not. It's hard to know whether the subjects' behavior on the test is truly reflecting *perception*. Possibly what's happening is something like this: The test questions are very hard to answer with any certainty, based on perceptual information alone. Maybe when confronted with a trio of very similar colors, English speakers tell themselves, "Hmm, these all look so close to each other, it's hard to tell which one of them is slightly more different from the other two. Still, two of these would probably be called 'green' and one of them would probably be called 'blue,' so I'm going to guess that the 'blue' one is the odd one out." In this case, the availability of different labels may offer a handy strategy to resolve a tricky perceptual task, rather than directly affecting how the colors themselves are perceived.

What would it take to convince you that language truly does interfere with perception? Perhaps if we design a study in which subjects are made to respond to the perceptual test as quickly as possible, this would make it hard for them to invoke a fallback strategy based on color names. In fact, we could then use the response times themselves as evidence for whether color distinctions are easier to make across word boundaries.

This was exactly the strategy used by Jonathan Winawer and his colleagues (2007) in exploiting the differences in color vocabulary between English and Russian. English speakers are able to use a single word (*blue*) to describe all the shades of blue pictured in **Figure 12.10A**, but Russian speakers are required to be more precise and distinguish between light blue (*goluboy*) and dark blue (*siniy*). (And, as it turns out, Russian speakers make the cut between *siniy* and *goluboy* at about the same place on the blue spectrum as English speakers who are forced to decide whether a color is light blue or dark blue.) Winawer and his colleagues devised a simple timed task in which English- and Russian-speaking subjects saw three squares of color, with one square on top and two on the bottom, as in **Figure 12.10B**. In each trial, one of the bottom squares was identical in color to the one on the top; subjects were told to quickly press a button on

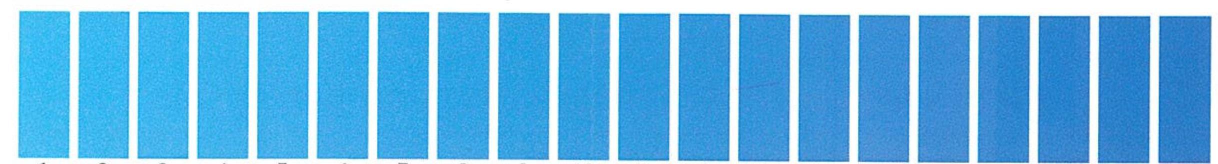

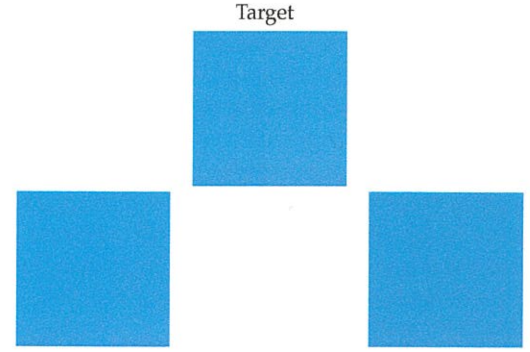

Figure 12.10 (A) Stimuli from the color perception study by Winawer et al. (2007). (A) All of these stimuli can be referred to by the English label *blue*. Russian speakers, however, divide the set up into *goluboy* (light blue) and *siniy* (dark blue), with the boundary between them typically occurring between stimuli 8 and 9. (B) An example trial from the study. To respond correctly, subjects had to press a button on the left to indicate that the lower left square matched the top square in color. (Adapted from Winawer et al., 2007.)

their left if the left square matched the top square, and a button on their right if the right-hand square was the correct match.

Not surprisingly, it took both groups of subjects longer to respond if the two bottom squares were very close in color than if they were further apart. But the Russian speakers performed a bit differently than the English speakers did: if a trial contained two colors that were very similar but sat on opposite sides of the *siniy/goluboy* fence, the subjects' responses were faster than if the two colors would both be classified as either *siniy* or *goluboy*. The existence of the two distinct words appeared to sensitize them to that particular distinction in color. English speakers, on the other hand, showed no advantage for colors that straddled the *dark blue / light blue* divide.

Evidence from reaction times is a fair bit more compelling than the judgment task used by Kay and Kempton. You may still not be fully convinced, though, that people's behavior in this task reflected an involuntary, *automatic* response to perceiving color rather than the use of verbal information to make a *decision* about color. The skeptics among you are invited to work through the details of the intriguing ERP study summarized in **Box 12.4**. In this study, effects of color vocabulary for English and Greek speakers were apparent within about 100 milliseconds of the presentation of colored stimuli. Thus, it's beginning to look like perhaps language truly does influence color perception.

12.4 Adjusting the Language Dial

How to silence the Whorf effect

There's a fascinating twist to the Whorfian story. In the study by Jonathan Winawer and his colleagues, there was one condition in which the Russian speakers performed just like the English speakers in that they were no better at detecting differences between stimuli that fell into different *siniy* versus *goluboy* categories than they were at noticing within-category differences. This happened when the subjects were saddled with an extra memory task: just before seeing each color trio, they read an eight-digit number that they knew they

BOX 12.4

ERP evidence for language effects on perception

An ERP study by Guillaume Thierry and his colleagues (2009) suggests that language influences on color perception can occur quickly and automatically, even in an experimental task in which subjects are not told to focus on color at all. In this experiment, EEG recordings were taken as English- or Greek-speaking subjects saw colored shapes, most of which were circles. Following instructions, the subjects monitored for the occasional square by pressing a button when they saw one. The squares, however, were simply decoys—the true stimuli of interest were circles that deviated in color from the majority of the other circles in the trial (**Figure 12.11**). The researchers anticipated that a circle of an unexpected color would trigger a visual mismatch negativity effect (vMMN), which is an ERP component that reflects the automatic detection of a change in the visual stimulus. It's generally assumed that the vMMN is a *preattentive* effect, that is, it occurs regardless of whether people are paying conscious attention to the aspect of the stimulus that changes, as a sort of unconscious registering of surprise.

Figure 12.11 Examples of trials in the study by Thierry et al. (2009). For each of these four trials, subjects were instructed to press a button when they saw a square (the target). EEG data were recorded and analyzed for the subjects' responses to the deviant—circles that appeared in a different color than the majority of the items in that trial.

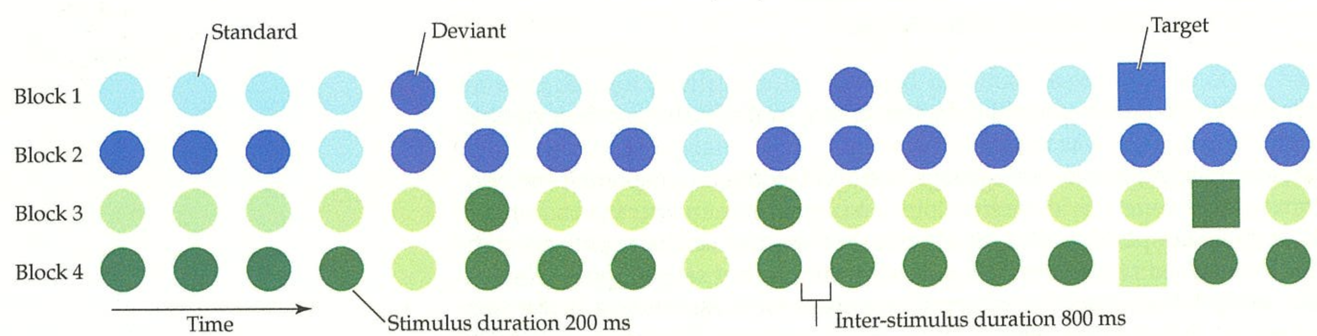

would have to correctly identify after responding to the color test. This meant that while performing the color test, they were mentally rehearsing the strings of digits and, presumably, clogging up their language system in the process. Somehow, blocking the possible activation of the color names eliminated the effects of color vocabulary on perception. And it really does seem to be that *verbal* interference was the culprit; when the Russian subjects had to remember spatial grid patterns instead of digit sequences, they once again showed heightened sensitivity to color differences across the *siniy/goluboy* boundary.

So, do color names alter the perception of color? Well, the answer seems to be yes, but apparently only if people can activate those linguistic labels while performing the perceptual task. This conclusion is supported by an inventive study carried out by Aubrey Gilbert and colleagues (2006), in which the researchers tackled the question from a very different angle.

As you'll remember from Chapter 3, words tend to be activated mainly in the left hemisphere of the brain—some evidence for this came from the classic experiments with split-brain patients that I described earlier. In those studies, patients whose hemispheres had been surgically disconnected were able to identify objects but not name them if the pictures of the objects appeared in the left visual field. Moreover, even in people whose hemispheres were properly connected, studies showed that words were recognized more efficiently when they were presented to sensory organs on the right side of the body, which are directly connected to the left hemisphere—for instance, in dichotic listening

BOX 12.4 *(continued)*

The critical question was whether the subjects' native language would affect the vMMN. As shown in Figure 12.11, the circles were either green or blue, and varied within trials as to whether they were light or dark shades of those two colors. In English, the same word *green* or *blue* can be used to apply to both light and dark versions of these colors, but Greek differentiates between light blue (*ghalazio*) and dark blue (*ble*), while using a single name (*prasino*) for light green and dark green. Hence, for Greek speakers, the perceived difference between the light and dark blues was expected to be greater than for light and dark greens, thereby eliciting a larger mismatch effect. English speakers, on the other hand, should show no difference in the size of the mismatch effect. This hypothesis is supported by the ERP results of the study (**Figure 12.12**).

Figure 12.12 ERPs elicited by the standard and deviant circles, summarized over numerous recording sites. (A) Results for native English speakers. (B) Results for native Greek speakers. As predicted, the native Greek speakers showed a larger difference in brainwave activity at around 200 ms for standard versus deviant blues as compared with the difference for standard versus deviant greens. For native English speakers, the mismatch effect was the same for greens and blues. (Adapted from Thierry et al., 2009.)

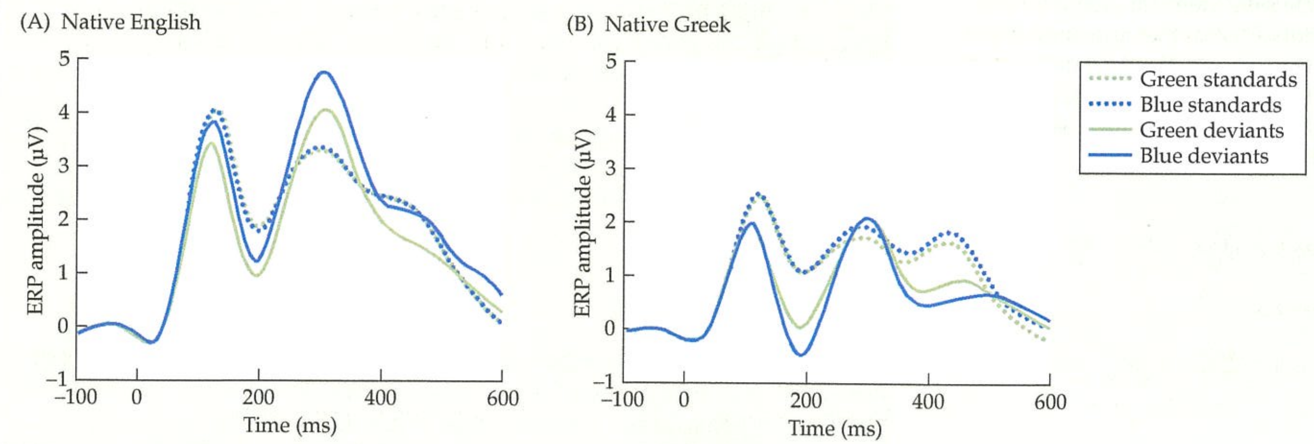

tests, where subjects heard a different word in each ear, they had an easier time recognizing the words presented to the right ear.

Given that word representations are more active in the left hemisphere, Gilbert and his colleagues wondered whether a person's color vocabulary would influence perception differently depending on whether the stimuli were presented to the left or the right visual field. They devised a study in which subjects saw a ring of colored squares, with one "oddball" square of a slightly different hue than the others. As in the previous studies, the objective differences between the color squares were kept the same for all of the trials, but the squares could represent either two different linguistic categories (green and a nearby blue) or just one (two shades of blue, or two shades of green), as shown in **Figure 12.13**. The subjects' job was to quickly press one of two buttons to indicate whether the odd-colored square appeared on the left or right side.

When the target oddball square appeared in the subjects' right visual field (with the visual information being processed in the verbal left hemisphere), people were relatively fast at distinguishing it from the other squares if it fell into a different linguistic category (for example, a blue target square among a set of green ones); they were slower to respond if the oddball bore the same label (for example, a blue among slightly different blues). This showed, once again, that language can enhance sensitivity to certain subtle color differences. But when the oddball square appeared in the *left* visual field (with the visual in-

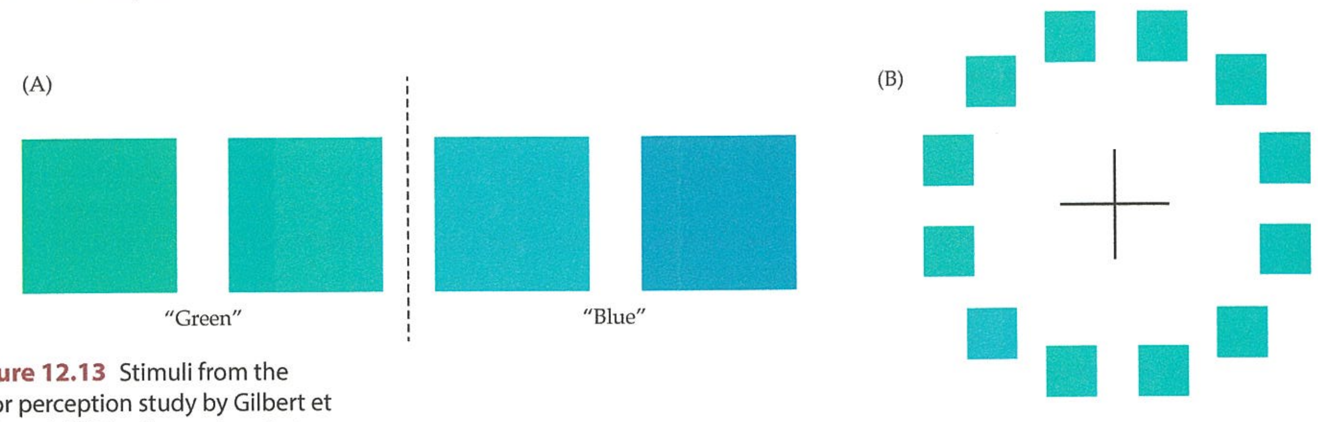

Figure 12.13 Stimuli from the color perception study by Gilbert et al. (2006). (A) The four colors that were used as stimuli, and the boundary that separates the greens from the blues. (B) The arrangement of stimuli in an example trial. All color squares in the circle were identical except for one "oddball" color that appeared either to the left or to the right of the cross in the middle, where subjects were instructed to focus their eyes. (From Gilbert et al., 2006.)

formation processed in the right hemisphere), the language effect evaporated, with no heightened sensitivity for different-named colors.

Results like these argue against a version of the Whorf hypothesis in which language has permanently altered the perception of color. Instead, language imposes itself on perception rather selectively—it appears that certain purely perceptual categories can exist and continue to operate outside of language's sphere of influence. But when the corresponding linguistic categories are highly active, they can play an important role in our perceptual experiences. You might think of language not so much as the teacher or guide of perception, but as an opinionated and vocal consultant. Under some circumstances, its opinions are muffled and perception carries on alone without the influence of language.

Beyond color words

Much of the Whorfian debate has been fought in the arena of color perception, and the abundance of experimental work on color has certainly helped to sharpen ideas about how language might influence perception or thought. But how much do the findings about color tell us about the relationship of language and thought more generally? Color perception by its nature involves making subtle distinctions about gradations of hue or brightness. But in other conceptual domains, categories might be much more sharply defined—think of giraffes and elephants, for example. In these cases, language might play a less important role in influencing judgments about categories. On the flip side, in more abstract conceptual domains, some categories could be more difficult to think about at all without the help of language—for instance, try to think *non*-linguistically of concepts such as a week, democracy, theory, or a contract.

Recent research suggests that the findings from experiments with color extend to at least some other concepts. In one study, Aubrey Gilbert and colleagues (2008) applied the same methods they'd used for studying color to look at concepts involving animal categories like cats and dogs. They used the same visual search task in which subjects were asked to identify one image that differed from others arranged in a ring (**Figure 12.14**). Just as in the color study, the images could come from different categories (an image of a cat appearing in a ring of dog images) or from the same category (two different images of cats). In general, subjects were faster to identify the oddball if it came from a different category than the surrounding images. But this cross-category advantage was greater if the oddball appeared in the right visual field than if it appeared in the left. This pattern makes sense if the rapid access of the *names* for cats and dogs in the left hemisphere allowed people to be faster at detecting the visual difference between the oddball item and its neighbors. In an intriguing variant of the experiment, the researchers tested a split-brain patient, who

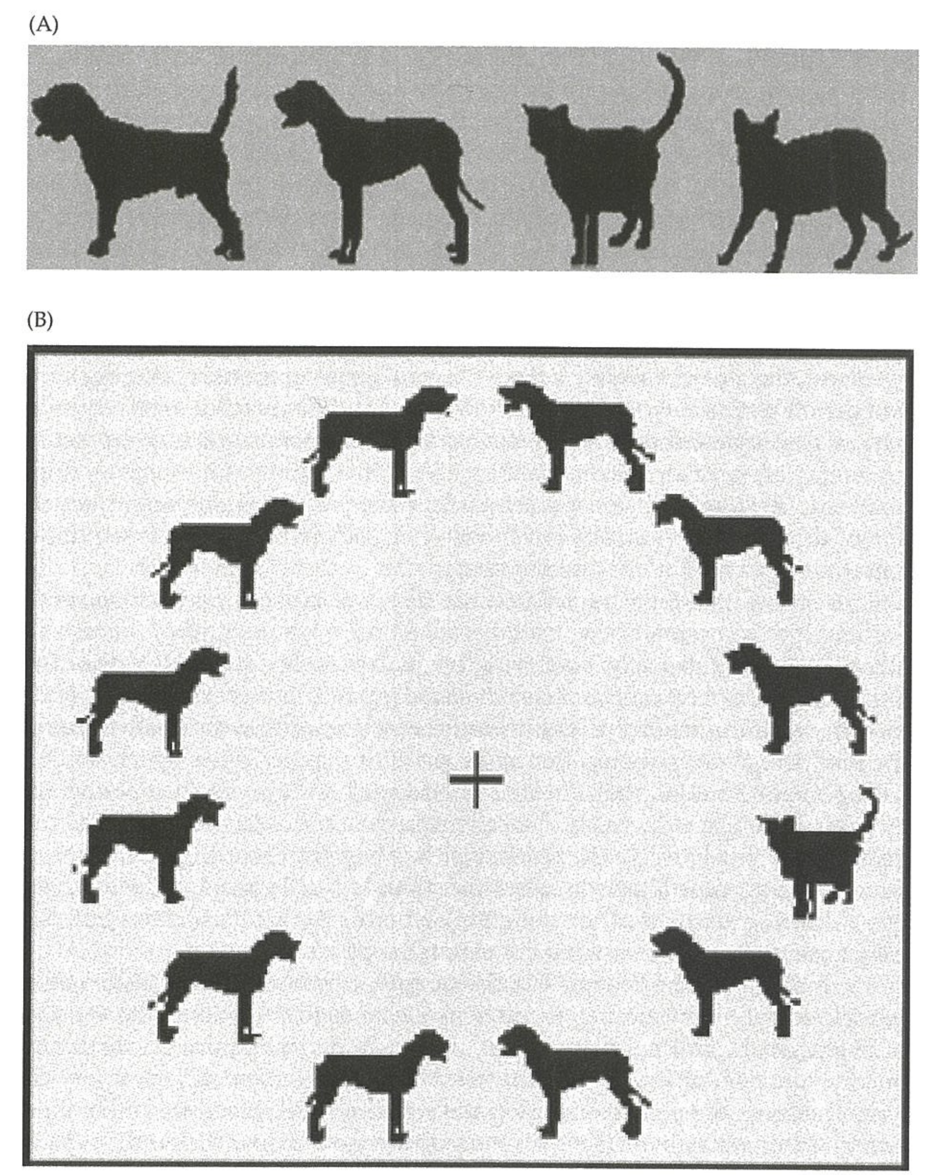

Figure 12.14 (A) The stimuli used in Gilbert's experiment (2008), including two different images in each of the cat and dog categories. (B) An example trial, illustrating an oddball that is from a different category than the surrounding images. (From Gilbert et al., 2008.)

showed an even more dramatic difference between the left-field and right-field results. In her case, the cross-category advantage disappeared entirely when the oddball images were presented in the left visual field. This suggests that when the two hemispheres weren't able to communicate with each other (that is, when language was prevented from exerting an influence on the task in the left-field condition), it was no easier to spot the difference between cat and dog images than it was to distinguish one cat image from another. So, much like the experiments with color terms, this study shows that conceptual categories can operate independently of linguistic categories, but that in some circumstances, the two become yoked to each other.

What about even more complex concepts—for instance, concepts that involve *events* rather than just objects? Not surprisingly, it's easy to find examples of languages that differ sharply in how they encode a complex event. For example, there are interesting crosslinguistic differences in the information that gets packed into verbs of motion. In English, we have a copious assortment of verbs to describe how a person might move his body from one location to another—