Trends in Cognitive Sciences

CellPress REVIEWS

Review

Effects of Language on Visual Perception

Gary Lupyan,[1,*] Rasha Abdel Rahman,[2] Lera Boroditsky,[3] and Andy Clark[4,5]

Does language change what we perceive? Does speaking different languages cause us to perceive things differently? We review the behavioral and electrophysiological evidence for the influence of language on perception, with an emphasis on the visual modality. Effects of language on perception can be observed both in higher-level processes such as recognition and in lower-level processes such as discrimination and detection. A consistent finding is that language causes us to perceive in a more categorical way. Rather than being fringe or exotic, as they are sometimes portrayed, we discuss how effects of language on perception naturally arise from the interactive and predictive nature of perception.

Highlights

Our ability to detect, discriminate, and recognize perceptual stimuli is influenced both by their physical features and our prior experiences.

One potent prior experience is language. How might learning a language affect perception?

We review evidence of linguistic effects on perception, focusing on the effects of language on visual recognition, discrimination, and detection.

Language exerts both off-line and on-line effects on visual processing; these effects naturally emerge from taking a predictive processing approach to perception.

> Even comparatively simple acts of perception are very much more at the mercy of the social patterns called words than we might suppose [1].

> No matter how influential language might be, it would seem preposterous to a physiologist that it could reach down into the retina and rewire the ganglion cells [2].

Language as a Form of Experience That Affects Perception

What factors influence how we perceive the world? For example, what makes it possible to recognize the object in Figure 1A? Or to locate the 'target' in Figure 1B? Where is the head of the bird in Figure 1C? Why do we perceive some colors in Figure 1D as more similar than others? Research on visual perception has sought to answer such questions by focusing largely on the physical properties of the stimuli and their interactions, (e.g., [3–5]). However, it has been long appreciated that what we perceive is determined not only by the physical properties of the current input, but also by our perceptual history. For example, consider how much harder it is to read upside-down text [6] or to match a stranger's versus a friend's photograph to their actual face [7]. Such effects of prior experience on perception can be observed not only for arguably 'higher-level' processes such as recognition, but also for 'lower-level' processes such as amodal completion [8], computing shape from motion [9], and computing 3D structure from object contours [10] and binocular disparity [11]. Although there is continued debate on the 'modularity' of some of these processes [3] (cf. [12,13]), there is relative consensus that what we perceive is importantly shaped by prior experience, (e.g., [14]). But what kinds of experiences matter?

A growing number of studies show that perception is affected by language. Most uncontroversially, experience with language affects perception of language. In learning English, we learn to perceive certain speech sounds as being functionally similar. This process of categorization distorts our perception, causing us to perceive physically equidistant sounds as more or less similar depending on our linguistic experience [15,16]. Such effects do not end at the level of individual speech sounds. Our experience with grouping certain combinations of speech sounds into larger units (such as words) causes us to perceive the same sounds differently, depending on which word they are embedded in [17,18]. Learning to read has profound impacts on a large part of our visual cortex (e.g., [19]); the consequences can be readily appreciated by comparing the experience of looking at a familiar versus an unfamiliar writing system.

[1]University of Wisconsin-Madison, Madison, WI, USA
[2]Humboldt-Universität zu Berlin, Berlin, Germany
[3]University of California, San Diego, CA, USA
[4]University of Sussex, Brighton, UK
[5]Macquarie University, Sydney, Australia

*Correspondence
lupyan@wisc.edu (G. Lupyan).

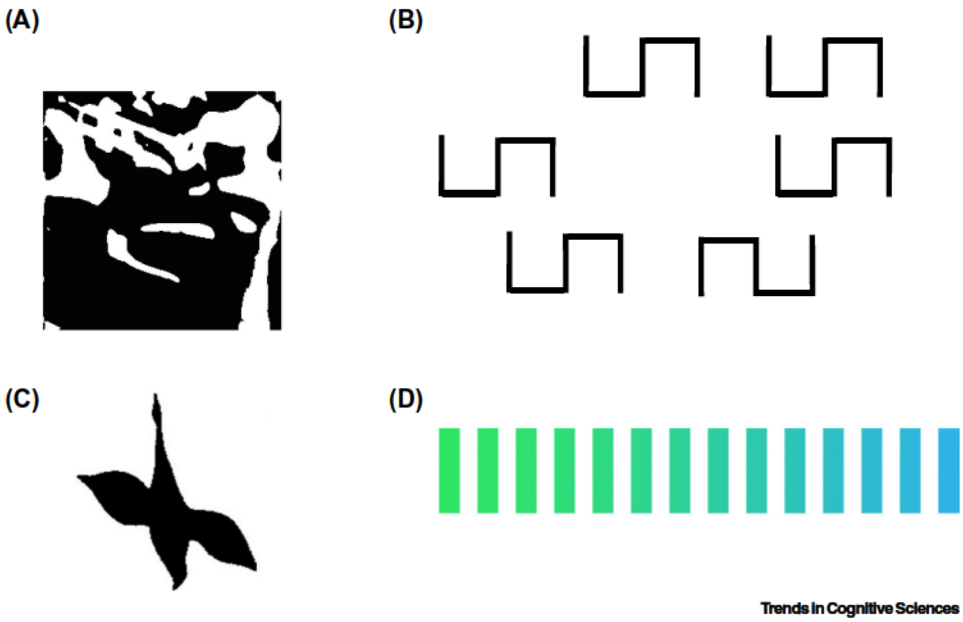

Figure 1. Some Examples of Stimuli Used to Study Influences of Language on Visual Perception. (A) A two-tone 'Mooney image' that becomes much easier to recognize after a verbal hint. (B) A sample visual search display. The target (odd-one-out) shape becomes easier to find when informed that the shapes are rotated numbers. (C) The location of this bird's head depends on expectations set up by reading descriptions of motion. (D) Psychophysically equidistant colors that may become perceived more categorically under the influence of the names 'green' and 'blue'. See text for details.

More controversially, experience with language and its in-the-moment use can affect perception of nonlinguistic material. In the next sections we review evidence of how language affects different aspects of perception. We focus on visual perception, not because effects of language on perception are limited to the visual modality, but because this is where most work has concentrated. We begin by describing two types of linguistic effects on perception: off-line effects in which long-term experience with a specific language affects how people subsequently experience certain perceptual inputs, and on-line effects in which some aspect of language, such as an object's name, interacts with in-the-moment visual processing. We then review empirical evidence of linguistic influences on perception, distinguishing effects of language on (visual) recognition, discrimination, and detection. We discuss whether differences in linguistic experience give rise to differences in perceptual experience and consider how the variety of the findings we review may arise from the workings of a perceptual system that seeks to minimize prediction error [20–22].

Effects of Language on Recognition, Discrimination, and Detection

Effects of Language on Recognition

Seeing a chair and recognizing it as such requires relating the current perceptual input to a previous state (this thing before me is a chair because it looks more like things I've seen previously that are chairs compared with things that are not chairs). To recognize is to categorize [23]. Might the categories we learn as part of learning a language affect visual recognition? Because we are so well-practiced in recognizing conventionally presented images, a productive method of studying recognition involves presenting incomplete or otherwise fragmented images. For example, only about 25% of people can recognize the object shown in the 'Mooney' image in Figure 1A. That is, given this input, most people's perceptual systems fail to make sense of it. The same input can be made meaningful if we are allowed to glimpse a conventional version of the image of the sort we have more experience recognizing. Glimpsing the conventional image changes not just the ultimate recognition success, but causes measurable changes at various levels of the visual hierarchy [24,25], for example, affecting people's ability to detect image contours, a relatively low-level perceptual task [12].

Could similar increases in recognition be achieved with the aid of language? In a now classic paper arguing for the cognitively impenetrable nature of visual perception, Pylyshyn argued that 'verbal hints [have] little effect on recognizing fragmented figures' [26] such as those in Figure 1A. Contradicting this assertion, a recent study [13] found that recognition can be substantially increased by simply giving verbal hints. For example, providing people with 15 alternatives as choices increases correct recognition of Figure 1A to 89%. Recognition can be increased to nearly the same level by providing a superordinate cue such as 'musical instrument'. One way to think about such effects is that the labels help to form hypotheses against which the incoming visual input can be tested. An otherwise uninterpretable perceptual input becomes interpretable when processed in light of hypotheses generated by language (see the section Making Sense of the Evidence).

Can language change our recognition of a visual scene, even when the visual scene is unrelated to the linguistic content? In one study [27], people viewed an image of an ambiguously facing bird (Figure 1C) after they had seen real upward or downward motion or read a linguistic description of physical motion (entirely unrelated to birds). Afterwards, participants were simply asked to draw a worm in the bird's beak. People's recognition of the bird's head was affected similarly by viewing real motion and by reading stories describing physical motion.

It may be tempting to interpret some of these effects as effects of language on downstream interpretations of perceptual processes which themselves remain free from linguistic influence [28]. One way of distinguishing processes that are more clearly perceptual from those that may reflect higher-level semantic or downstream decision-making processes is by using event-related brain potentials (ERPs). Some ERPs are associated with the processing of low-level visual properties as lightness and contrast (P1 component, peaking at around 100 ms [29]) or high-level perception of objects and faces (N1 component, peaking at around 150 ms) [30,31]. Other components such as the N400 are better described as tracking more amodal (or multimodal) semantic processing [32,33]. These studies find that holding perceptual experience constant while varying how much one learns about a novel object, through language, affects early visual processing of objects. Learning an object's name or function decreases subsequent P1 amplitudes during passive viewing [34,35]. The early changes in visual processing indexed by these electrophysiological markers are associated not only with changes to visual recognition [36] but also in changes in discrimination and conscious awareness, as the next two sections describe.

Effects of Language on Discrimination

While recognizing an image requires that it be discriminated from possible alternatives, it is certainly possible to successfully discriminate images without recognizing them. For example, we can discriminate between the letters of a novel alphabet without being able to recognize the letters. We can use discrimination tasks to tell whether effects of language on perception extend beyond recognition.

Many investigations of how language affects discrimination have been done in the domain of color. Indeed, the finding that color names may influence color discrimination is often viewed as a key test of the Whorfian hypothesis [39–41], even though Whorf never proposed that language should affect color perception. This topic has been extensively reviewed elsewhere [41–45], but we highlight a few findings that show the variety of ways that language affects color perception.

What makes color an excellent domain of study is that the physical stimulus is continuous, but we tend to talk about it categorically. Learning color words necessarily requires learning to group colors into categories named by our language, treating different colors that have the same name as more similar than psychophysically equidistant colors named with distinct labels. This in turn may cause us to perceive color in a more categorical way. Adding to the intrigue, some languages have many color words, while others have few or none (Box 1).

Box 1. Cross-Linguistic Differences in the Language of Perception

Although all languages provide ways for describing perceptual experiences, and there are broad similarities in the naming systems, there are also surprisingly large differences in the details. Take just about any word that describes a perceptual property ('spicy', 'triangle', 'loud', 'blue') and you will find none to be universal [121,122]. These lexical differences have often been minimized by pointing out that people can always invent words if they are needed [123]. But this confuses the potential for the actual. It is true that over decades or centuries languages change to meet new communicative demands. But this does not mean that an individual speaking a language lacking color words can, on their own, invent them! In the course of learning English, children become experts at naming (categorizing) some part of color space. In comparison, Lao speakers learn to categorize a far smaller part of color space (Figure IB). But, compared with English speakers, they have more consistent classification of certain odors [124]. People speaking either language can become expert categorizers in any domain, but having the requisite words as part of the core vocabulary ensures that all speakers do.

Another point of confusion concerns the relationship between the number of terms in a language and their use. It is not that languages with fewer color terms necessarily have terms that apply to larger regions of color space (see [125] for a model that connects naming consistency to communicative efficiency]). Rather, languages with fewer color terms sometimes have large regions of color space without conventional names [124,126] (Figure IB). Without a conventional label, it is difficult (perhaps impossible) to cue the category. How do you get someone to attend to blues without being able to use 'blue'?

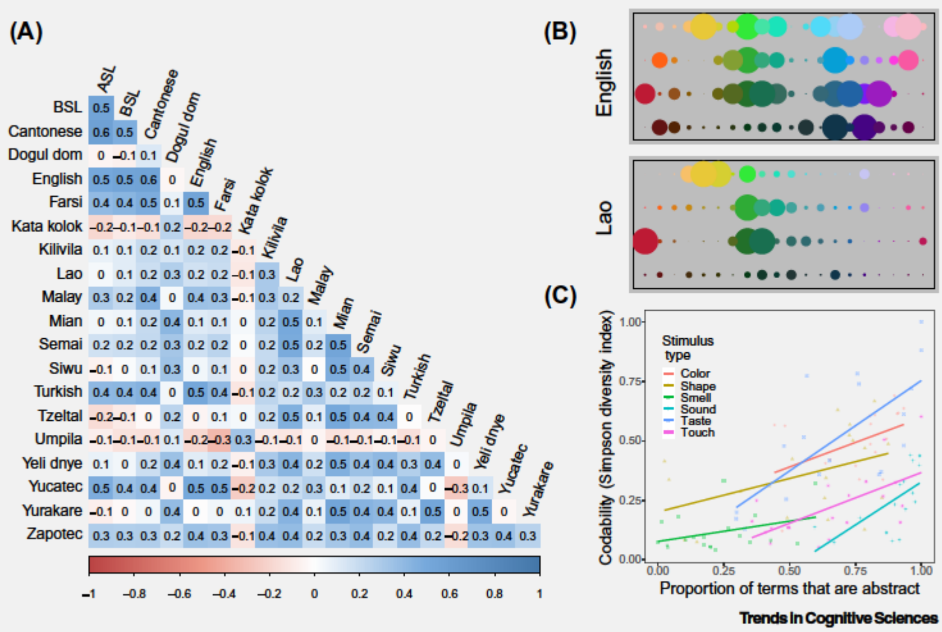

Figure I. Cross-Linguistic Differences in the Language of Perception. (A) Cross-linguistic correlations in color nameability for 40 Munsell color chips of maximum saturation; smaller numbers indicate that the most nameable colors in one language are not the same as those in another. (B) Nameability of maximally saturated colors in English and Lao (colors shown are RGB approximations of Munsell color chips). (C) For all perceptual modalities, greater nameability is associated with greater use of abstract terms (e.g., 'yellow', 'scratchy') as opposed to source-based terms ('lemon', 'sandpaper'). Data shown are based on reanalysis of [124] and are made conservative by including phrasal heads only; 'dark blue' and 'light blue' are counted as the same response. Abbreviations: ASL, American Sign Language; BSL, British Sign Language RGB, red/green/blue.

Language has been shown to influence color perception in two main ways, roughly in line with so-called off-line and on-line influences schematized in Figure 2. First, are studies showing that people speaking different languages show somewhat different patterns of color discrimination. Russian speakers who obligatorily distinguish between light and dark blues in language show a relative advantage in discriminating colors that span the light/dark blue boundary compared with English speakers [46]. Similar differences in discrimination have been found when comparing speakers of other languages that differ in their color naming such as English and Korean (e.g., [47]). Use of electroencephalography (EEG) has allowed researchers to see that cross-linguistic differences in color discrimination tasks emerge as early as 100 ms after viewing a color [48–50]. Similar effects on early visual processing are observed processing objects that are lexically distinguished in one language, but not another [51].

Second, are studies showing that color discrimination can be altered within an individual by manipulating linguistic factors. Having adults learn new color categories (analogous in some ways to children learning to produce and comprehend color labels) can induce more categorical color perception [52–54], a finding consistent with both off-line and on-line effects of language (Figure 2A). At the same time, several lines of evidence show that interfering with labels can reduce or eliminate categorical effects [55] and negate cross-linguistic differences [46]. Conversely, overt use of color labels during the task has been found to exaggerate categorical influences on perception. For example, presenting participants with a color word (e.g., 'blue') immediately prior to requiring them to discriminate one color from three others (a simultaneous odd-one-out discrimination task) dramatically improves their ability to distinguish blues from greens while slightly reducing accuracy of distinguishing highly typical blues from slightly less typical blues [38]. Incorporating category typicality (as differential association between an exemplar and its label) into the model makes a further counterintuitive prediction: as the activation of a category label shifts perceptual representations toward category prototypes, the shift will not only impact colors straddling the category boundary, but also spread to colors around the middle of the category, as schematized in Figure 2B. Indeed, hearing a color label such as 'blue' was found to increase the accuracy with which people distinguished typical blues from less typical blues [38]. Another intriguing effect is that of stimulus order. Imagine seeing a green color swatch G_1. Then after a

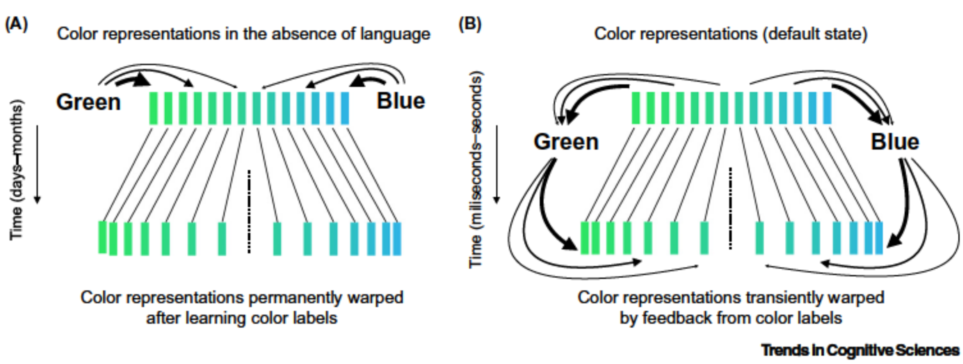

Figure 2. Schematic of (A) Off-Line and (B) On-Line Effects of Color Labels on Color Representations. (A) Colors that are originally represented as equally discriminable may become gradually warped by the categorization required for comprehending and producing color names. (B) Categorical effects may be produced by a color percept activating a verbal label which then feeds back and warps the color representation into a more categorical form in-the-moment. Note that because the labels 'blue' and 'green' are more strongly associated with some blues and greens than others (denoted by line thickness), a more categorical representation leads to expansion not only around the category boundary, but also around the middle of the category (see also Figure 2 in [37]) leading to improved discrimination of atypical members from slightly more typical ones [38]. On-line effects may be caused by covert labeling (automatic activation of labels by perceptual inputs) and further exaggerated by overt labeling such as actively naming a color or reading/hearing a color term.

short delay, you are asked to select whether you saw G_1 or G_2. It turns out that accuracy is substantially better if G_1 is the more typical green than if G_2 is the more typical green [56,57]. One explanation is that the category label is activated more strongly by the more typical color and the label then feeds back to warp the representational distance between the two color stimuli, enabling more accurate matching [38]. (The effect of presenting the most prototypical stimulus first is predicted to reverse if both G_1 and G_2 are near the category prototype.)

Somewhat more circumstantial evidence for the involvement of language in color perception comes from findings that in color discrimination tasks people show a larger between-category advantage in the right visual field (which projects to the left hemisphere) than in the left visual field [47,53,55,58,59], that verbal interference selectively affects between-category discrimination in the right visual field [55], and that the lateralization difference is already observed in early visual processing as measured by EEGs [60]. In related studies using functional neuroimaging, it has been found that color discrimination tasks evoke activity in cortical regions associated with naming (e.g., the left middle superior temporal gyrus), that this activation is stronger for colors viewed in the right visual field [61], and that discriminating easy to name colors evoked greater activity in 'naming' regions compared with discriminating harder to name colors [62]. Not everyone agrees that finding categorical perception in the left hemisphere implies that language plays or has played a causal role. One alternative is that the observed lateralization effect may stem from the left hemisphere being more specialized for categorical processing regardless of language [63,64]. The lateralization effect has also not always replicated [65,66], suggesting that it may be more fragile than originally thought, a critique that has also been levied against studies showing effects of labels on color memory (cf. [67,68]) (see Outstanding Questions).

Effects of language on visual discrimination are not limited to colors. Ascribing meaning to simple objects such as referring to the shapes in Figure 1B as 'rotated numbers', leads to more efficient visual search [69]. Search can likewise be improved simply by referring to an object by its name. For example, when repeatedly searching for a 'B' among 'p's, hearing 'find the B' shortly before the search display speeds reaction times compared with hearing 'find the target', even though the label was uninformative because 'B' was always the target [70]. People's visual discrimination of novel 'alien' creatures is improved by learning to associate them with dissimilar verbal labels (one is 'loud, nocturnal, strong' and another is 'sticky, soft, wet') compared with semantically similar labels ('loud, nocturnal, strong' versus 'loud, heavy, strong') [71]. The earlier-discussed study with Mooney images [13] showed not only that verbal hints improved recognition, but also found that after being exposed to verbally labeled Mooney images, participants were better able to visually discriminate one Mooney image from another compared with a condition in which the same images were viewed equally often, but without being labeled. This effect was mediated by changes in P1 amplitudes and increases in alpha-band oscillations, especially in the left hemisphere. As we discuss in the section 'Making Sense of the Evidence', these results are consistent with labels preferentially activating category-diagnostic features, properties that most effectively distinguish category members from non-members, which results in better recognition of category members and more efficient discrimination between members and nonmembers [36,72].

Effects of Language on Detection

An even more basic perceptual process is simple detection wherein people are asked to indicate whether they have seen anything during a specific time period. Simple detection tasks require neither recognition nor discrimination (except for discriminating an object from the background). When shown briefly presented backward masked letters, people's ability to correctly indicate whether a letter (any letter) was present, was affected by hearing the letter's name immediately prior [73]. It is as though the name prepared the visual system to detect (not just recognize) a class of shapes denoted by the label. A shortcoming of this study is that although brief masked presentation makes stimuli difficult to see, it does not block semantic processing [74], leaving open the possibility that hearing a label affected downstream processes rather than something as basic as visual detection. A later study [75] used continuous flash suppression (CFS) to suppress perceptual processing at a much earlier level [76]. In a typical CFS experiment, a meaningful image is presented to one eye while high contrast flashing patterns are presented to the other. The flashing patterns act to suppress awareness of the image through a form of binocular rivalry. Hearing a word (e.g., 'pumpkin') prior to viewing a CFS display made an otherwise invisible image (e.g., of a pumpkin) visible. The detection advantage was limited to trials on which the verbal cue matched the suppressed stimulus: hearing 'pumpkin' did not cause people to confabulate seeing pumpkins where none existed. Because images suppressed through CFS are not processed semantically, this detection benefit cannot be ascribed to downstream effects of language. Subsequent work supported the interpretation that the detection advantage stems from labels activating the shape of the to-be-detected object [77].

Another method of investigating the influence of language on visual awareness involves detecting stimuli in a stream of rapidly presented images. A classic finding is that when participants are asked to detect two targets (T_1 and T_2) within a visual stream of non-target images, presenting T_1 causes people to miss T_2 if T_2 occurs 200–500 ms after T_1, a so-called attentional blink [78]. A popular explanation is that ongoing processing of T_1 hinders T_2's access to a second processing stage that is necessary to produce a durable representation of the stimulus [79]. In a recent study [80], researchers used pictures of rare objects as the T_2. When subjects previously associated the pictures with verbally described functions (e.g., 'this is an incubator for chicken eggs'), they showed an increase in conscious detection. This effect was predicted from modulations of the P1 component approximately 100 ms after stimulus presentation. In another study using the attentional blink paradigm it was shown that native Greek and Russian speakers, who distinguish categorically between light and dark shades of blue, showed boosted detection of T_2 when it was marked by verbal contrast [48]. A P1 modulation for within- and between-category colors was registered in Greek participants (Russian speakers were not tested), and this modulation predicted their behavioral advantage on trials with lexically discriminated colors. By contrast, German speakers who do not habitually refer to light and dark shades of blue by different names showed no behavioral or electrophysiological differences between blue and green targets. In at least some cases, our native language predicts what we will consciously perceive.

Making Sense of the Evidence: Predictive Processing and the Reach of Language into Perception

On a naive view of perception, the idea that language can affect what we see is absurd. How can language, a high-level cognitive process unique to humans, affect the seemingly low-level mechanisms subserving visual processing? In the words of the opening quote, language cannot 'reach down into the retina and rewire the ganglion cells' [2]. But our conscious visual experiences and the processes involved in even the simplest of perceptual tasks cannot be reduced to the firing of retinal ganglion cells [45,81]. On the view of perception as a process of predictive inference, investigated under the somewhat coextensive banners of 'predictive processing', 'active inference', and 'hierarchical predictive coding', perceptual experiences arise at the meeting-point of generative model-based predictions and sensory stimulation (evidence). Percepts reflect 'best guesses ' of the world and these guesses are informed by prior knowledge, current sensory evidence, and context-varying estimations of their relative reliability ('precision') [82–84]. This process is not unique to visual perception and may be a useful framework for processing in all modalities, including the perception of pain (Box 2).

Box 2. The Effect of Language on the Perception of Pain

Just as the proximate cause of visual experience is photons entering the retina, the proximate cause of peripheral pain is the stimulation of various thermal, chemical, and mechanical receptors. But just as visual perception cannot be reduced to retinal stimulation, the experience of pain is not reducible to stimulation of these peripheral receptors. Of particular relevance to this review is the phenomenon of *placebo analgesia*, in which the same stimulus is experienced as either more or less painful depending on a person's expectations (e.g., [110–113]). The phenomenon of placebo analgesia is broadly consistent with the idea that perceived pain reflects both the properties of the external stimuli and the person's (certainty-weighed) expectations [114,115]. Having expectations of more or less pain does not simply bias people to respond to an external stimulus more or less strongly. Rather, these expectations modulate nearly the entire physiological profile of pain [113,116]. Some brain regions (e.g., posterior insula) continue to code the actual intensity of the stimulus somewhat independently of people's expectations, but this is ultimately irrelevant when we consider that our subjective experiences of pain tracks with physiological pain responses such as heart rate, skin conductance, and pupil size which are all modulated by our expectations [114]. Verbally induced expectations in particular ('I will now apply a cream that will help reduce the pain') play an especially important role [110,117–119].

It may be tempting to brush aside this example of language affecting perception as that of language simply being used as a tool to set up expectations. In principle, these expectations can be set up through other, nonlinguistic methods, for example, by learning to pair arbitrary shapes with a more or less painful stimulus [114]. But consider the difficulty of nonlinguistically conveying to a subject the idea that spreading this cream on their left arm will, in 15 minutes, reduce the amount of thermally induced pain in their arm [117]. Even if, in principle, these expectations can be set up nonlinguistically (e.g., by associating the cream with a nonplacebo analgesic) what language adds is fast, precise, and highly flexible deployment of these expectations [120].

There are two main ways in which language impinges on the predictive process, which roughly map onto the off-line and on-line distinction in Figure 2. The first is what happens during language learning. Here, labels can be viewed as supervisory signals (just as they are in supervised neural networks that have shown impressive successes in image recognition). In effect, linguistic labels act like artificial tasks, prompting the learner to actively seek ways to discriminate the labeled positive examples from the negative ones. If this ball is labeled 'red', what is it that made it 'red' and not 'green'? Exposure to linguistically labeled cases provides important information concerning the predictability of sensory patterns.

It is easy to see how labeled cases are critical for developing domain expertise. A trained botanist can learn to recognize thousands of plants and species by sight. In the absence of the teaching signals provided by language, there would need to be sufficient non-label-invoking tasks to drive the correct 'wedges' through a dense and confusing representational space. The same process is also at work in the learning of other categories we may take for granted. In learning to name colors, we not only learn where in color space our particular language places lexical boundaries, but we learn that color is a domain worth separating from other aspects of visual experience such as texture and vividness, an abstraction we might not make were it not for (some) languages demanding that we learn a color vocabulary. An important further source of perception-relevant knowledge gleaned from language may come from exposure to its distributional structure. By simply minimizing prediction error of the labels (and their co-occurrences) as they occur in language input, it is possible to learn a surprising amount of semantic structure, including relationships between perceptual features (Box 3).

Understanding this off-line role of language (Figure 2A) is critical to understanding how language influences perception. The usefulness of 'blue' or 'pumpkin' in guiding perception hinges on first learning the association between these words and their referents. It is logically possible that as one learns a language, habitual use of its terms gradually yet permanently reshapes perceptual representations. Following this early 'formative' period, however, language may no longer be actively involved in perception. However, findings that manipulating language during a perceptual task (e.g., through verbal interference and trial-by-trial verbal cueing), affects performance on a

Box 3. Visual Knowledge from Language Statistics

Associating some discriminable stimuli with common names and others with different names, a process of acquired similarity and distinctiveness [127], may be just one of the ways in which language 'trains' perception. As we have long known, languages are not just form-meaning pairings, but are coherent symbolic systems [128]. Word meanings cannot be dissociated from word use. The meaning of 'bedroom' derives in part from the existence of contrasting words ('bathroom', 'kitchen') and the contexts in which these words are used [129]. What this means is that co-occurrence statistics between words can act as a kind of echo of real-world linkages and causal relationships. A learning system whose sole experience is language can learn a surprising amount about the world by running a self-supervised learning algorithm on the incoming language stream (e.g., by trying to predict the next word and minimizing the prediction error based on the word that is actually observed) [130,131]).

As these learning algorithms are applied to larger amounts of text, we are realizing just how much structure people have 'pushed' into the co-occurrence statistics of languages. Consider the observation that congenitally blind people, who have no perceptual experience with color whatsoever, nevertheless know quite a bit about the color of various objects and the similarity structure of color space [132–135]. Recent investigations showed that it is possible to recover this type of information from distributional semantics alone (cf. [136,137]). The finding that the distributional structure of languages contains rich information about visual appearance does not tell us how much people rely on it for learning about what things look like, but it does raise the possibility that distributional semantics may be an important teaching signal for our perceptual systems [138].

range of perceptual tasks are difficult to square with a purely off-line account. Rather, language appears to modulate perception in the moment.

This brings us to the second way language influences perception: by providing a categorical expectation within which incoming perceptual input is processed. Imagine that you need to recognize a cow as quickly and accurately as possible. This can be achieved by preactivating the visual features that distinguish cows from non-cows. This suite of features is precisely what a categorical label like 'cow' is well designed to activate. And indeed, people are better able to recognize familiar images like that of a cow when it is preceded by its categorical label ('cow') than equally informative but less categorical nonverbal cues such as the sound of a cow mooing [72,85]. (Though there is some evidence that nonverbal auditory cues may be more effective in simple detection tasks [86].) The idea that labels elicit categorical expectations can be cast in terms of 'category-based attention' [70,87,88]. For example, we can tell someone to attend to vehicles or to faces or to colors (instructions that rapidly warp neural representations across the visual hierarchy [89,90]). As experimenters, we often take for granted our ability to use language to guide attention in this way, but consider the difficulty of placing a person into that attentional state without being able to rely on language. We can try to use alternatives: a picture of some vehicle, some face, a patch of color. But these are necessarily specific [91] and have limited power in cueing a categorical state [38,72].

Viewed in this way, what makes linguistic cues different from nonlinguistic ones is their patterns of associations. Any perceptual experience of a car or cow or the color green is quite specific; with language, by contrast, the same words are used to refer to a range of objects/events/relations; people learn that 'car', 'cow', and 'green' can be used to refer to any member of these categories. This makes labels ideally suited for setting up the sorts of categorical (and relational) expectations that are difficult to set up effectively using nonlinguistic means. That said, the underlying mechanisms by which linguistic and nonlinguistic cues set up expectations are likely to be much the same.

In short: learning a language allows us to use words and larger verbal constructions as a metabolically cheap and flexible means of modifying predictive cascades in ourselves and in others, altering what top-down information is brought in and how much influence it has at different levels of perceptual processing.

This predictive coding perspective may help to make sense of the limits on the effects of language on perception and the often-observed task sensitivity. For example, people who speak languages with different color naming schemes do not appear to differ in how small a color difference they are able to perceive (i.e., size of their just-noticeable difference) [54,92]. Within a language too, the size of a just-noticeable difference is not well correlated with the placement of lexicalized color categories [93] and categorical color perception itself is surprisingly task-dependent [94,95]. These findings have led some researchers to wonder whether such effects are limited to tasks that are somehow linguistic rather than 'truly' perceptual [96]. For example, some have argued that 'because labeling is not an inherent part of a visual process, ...we should not expect it to have a significant effect on visual appearance or discrimination ...the closer a color task is to language, the more likely it is that it will be affected by the terms in one's language' [97].

The predictive coding framework obviates the need to decide how linguistic or perceptual a given task is and instead poses the question of whether linguistic guidance helps to reduce prediction error [21]. In a task requiring discriminating small (largely within-category) differences in hue, linguistic guidance would be expected to do little. In a task requiring people to remember an item's exact position [98] or color [40,99], the finding that people's memory is affected by categories (both linguistic and nonlinguistic) can be usefully modeled by merging continuous perceptual representations with more categorical (discrete) conceptual/linguistic ones [40], with the original perceptual representation left intact. Yet we continue to see influences of categories on tasks where memory demands are minimal (e.g., simultaneous discrimination or reproduction of a currently visible stimulus such as matching one color to another [40,99]). One might think that in such tasks our decision mechanisms could simply draw on the earlier more continuous representation prior to it being 'contaminated' by categorical codes. Yet, as the data show, this is not always possible. On an alternative 'perceptual warping' account, the merging of continuous and categorical information happens through top-down feedback of categorical representations onto (the more continuous) perceptual representations. On the warping account, the reason labels can affect visual discrimination and reproduction is that the automatically activated labels are altering the lower-level perceptual representations themselves [38]. The level at which language has its effect is expected to depend on where in the perceptual hierarchy a change will reduce prediction error the most (cf. [100]). This is, of course, a rather vague claim and we recognize that much more rigorous work is required to precisely connect the idea of error minimization to patterns of behavior in experimental tasks. As a starting point, attempts to explain top-down effects on speech perception in terms of minimizing prediction error [101,102] offer a promising direction. Despite different conceptions of where precisely the merging of perceptual and categorical information is happening, what all these accounts have in common is that behavior in even low-level perceptual tasks is determined by more than bottom-up perceptual inputs. For language to change perception it hardly needs to 'reach down into the retina' [45].

Differences in Behavior versus Differences in Subjective Experience

Does finding that people speaking different languages perform differently on some perceptual tasks mean that learning different languages causes us to literally see the world in a different way? Do results showing that language affects objective behavior on perceptual tasks mean that language affects how the world appears? Some think the answer is a clear no. One critic sums up such effects as 'a passing flicker, that only painstaking experiment can reveal, in no way creating a different way of seeing the world' [103]. The experiments we review here indeed make painstaking attempts to rule out alternative explanations for the observed changes in behavior, attempting to demonstrate the perceptual nature of these effects. Many of these differences in accuracy and reaction times are likely not accompanied by substantive changes to subjective experience. This observation raises two questions.

What matters more: objective differences in behavior or subjective differences in experience? Imagine two participants (A and B) performing the same difficult discrimination task. Both inform us that they could not distinguish the items whatsoever. Participant A's data reveals chance-level performance. Participant B performs substantially better and well above chance. Do we conclude that A's and B's perception was identical? Or do we put our trust in the observed difference in objective behavior? It is, of course, interesting to find dissociations between subjective experience and objective behavior, but faced with choosing a measure, there is good reason to choose one based on objective behavior.

Would we know a difference in subjective experience if it existed? The collective shock of the Internet and vision scientists on discovering #theDress (https://en.wikipedia.org/wiki/The_dress), and the continued shocks whenever new stimuli of this sort are discovered: the jacket, the shoe, Yanny or Laurel, shows just how strongly we assume that the same perceptual input should produce the same perceptual experience. It is only by comparing notes that we can appreciate that sometimes it does not. How often such discrepancies occur and whether they can be predicted by differences in our experiences (language among them) is yet to be determined. Importantly, effects of prior experiences on subjective appearance are not willy-nilly. There is no reason to expect #theDress to appear differently to people speaking different languages. There is a theoretical reason to expect that it would appear differently to people who wake up early versus those who sleep in, and in fact, it does [104].

Some have wondered why top-down effects on perception seem to be so difficult to experience subjectively compared with, for example, visual illusions [105]. The compellingness of the best visual illusions depends on our ability to flexibly manipulate the visual stimulus. For example, the Adelson grid causes us to see two identically shaded squares as having different lightness [106]. What makes such illusions so compelling is that we can experience changes to our subjective experience by covering up the inducer, at which point the two squares start to look identical. The problem is that we cannot manipulate conceptual knowledge or prior experience with language in this way. We cannot pause being English speakers or switch on-and-off our knowledge that a face being shown is our mother's. As experimenters, we can manipulate top-down influences somewhat, for example, by testing bilinguals in different languages [45], by downregulating impacts of language through verbal interference or noninvasive neural stimulation, or by upregulating them through overt presentation of labels [107]. When these manipulations do lead to differences in behavior, we can infer that our normal experience must be, to some extent, influenced by the linguistic factor being manipulated. Sometimes even these subtle manipulations lead to differences in subjective experience [108] (see also Box 4).

Concluding Remarks

The opening quotes contrast two perspectives on the relationship between language and perception. There is some irony that Sapir, whose view was not informed by any of the empirical findings we discuss, is much more on the mark than Pinker, who relegates the idea that language affects thinking and perceiving to a time when 'scientists were in the dark about how thinking works or even how to study it' [2].

At one time, both language and perception were thought to be informationally encapsulated, their outputs feeding into domain-general cognitive systems, but not influencing one another [109]. The evidence we review here shows that language influences basic perceptual processing, affecting performance on tasks such as discrimination and detection, tasks that might seem to be wholly perceptual in nature. Far from a radical claim, this conclusion may naturally follow from viewing perception as an interactive process seeking to minimize prediction error. Becoming

Outstanding Questions

What is the relationship between language learning and the development of perceptual expertise? For example, are children who learn color and shape words earlier, better able to selectively attend to these properties or to flexibly switch between them? Does better naming lead to better and more robust visual recognition or deployment of attention?

What kinds of cross-linguistic differences lead to the most reliable differences in the perceptual experience of their speakers? Are there specific types of words or specific grammatical devices that are especially influential in their guiding of perception? For example, words denoting finer-grained distinctions may exert stronger influences than words denoting coarser-grained distinctions [146].

Are there systematic differences in perceptual experiences between people who experience inner speech to a greater versus lesser degree?

What task properties are most important for determining the level at which language influences perceptual processing?

Can we come up with critical experiments to determine whether the increase in categorical perception due to language stems from warping of lower-level perceptual representations or a merging of continuous perceptual representations with more categorical ones?

Box 4. Language and Mental Imagery

Conventional perceptual experiences are triggered by perceptual inputs from the outside world. Most of us are also able to entertain quasi-perceptual states in the form of mental imagery. Like conventional percepts, mental imagery can be triggered by external perceptual events. For example, seeing a friend's face may trigger a vivid mental image of the last time we saw her. Other times, what triggers visual imagery is language. When we want someone to experience the view from the hotel window without actually being there to see it, we do it with language. Most people also regularly use language directed at themselves, triggering mental images through inner speech [139]. As is now well known, imagery and perception engage overlapping neural circuits [140–142] and evoke similar representations of object categories [143].

Processing language that describes consistent visual motion in one direction has been shown to produce sufficiently vivid mental images to cause direction-selective motion adaptation in the visual system (i.e., cause a motion after-effect illusion). One study [144] tested for motion after-effects following explicit motion imagery and after processing language containing literal or metaphorical motion (without any instructions to imagine). The results demonstrate that language-evoked mental imagery produced direction-selective adaptation in the visual system.

It is interesting to consider what our visual imagery ability would be like in the absence of language altogether. Would our ability to form mental images be compromised if we had no language to help cue them? In a recent study [145], researchers varied the amount of (verbally transmitted) knowledge participants learned about novel objects, which they then either saw or had to imagine. The authors reasoned that greater semantic knowledge would lead to a stronger ability to reactivate the object's visual properties through top-down feedback. Semantic knowledge modulated early components of both object perception and visual imagery with similar P1 amplitude modulations in both tasks. Greater semantic knowledge also led to faster onset of visual imagery. These results show that in addition to influencing perceptual processing of actually presented objects, verbally transmitted knowledge can shape mental imagery.

a language user requires that we become expert at categorizing thousands of visual percepts into named categories. This experience equips us not only with the ability to efficiently communicate about our perceptual experiences, but to then use these words to flexibly deploy task-relevant hypotheses within which incoming perceptual information can be made more meaningful (see Outstanding Questions).

Acknowledgments

We thank Sean Roberts for help with Figure IB. Preparation of this manuscript was partially supported by NSF-PAC #1734260 to G.L., by Horizon 2020 European Union ERC Advanced Grant XSPECT -DLV-692739 to A.C., and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2002/1 'Science of Intelligence' - project number 390523135 and grant AB277/6 to R.A.R.

References

1. Sapir, E. (1929) The status of linguistics as a science. *Language* 5, 207–214
2. Pinker, S. (1994) *The Language Instinct*, Harper Collins
3. Keane, B.P. (2018) Contour interpolation: a case study in modularity of mind. *Cognition* 174, 1–18
4. Biederman, I. (1987) Recognition-by-components: a theory of human image understanding. *Psychol. Rev.* 94, 115–147
5. Treue, S. *et al.* (1991) Human perception of structure from motion. *Vis. Res.* 31, 59–75
6. Erlikhman, G. *et al.* (2017) On the legibility of mirror-reflected and rotated text. *Symmetry* 9, 28
7. Young, A.W. and Burton, A.M. (2018) Are we face experts? *Trends Cogn. Sci.* 22, 100–110
8. Vrins, S. *et al.* (2009) Bricks, butter, and slices of cucumber: investigating semantic influences in amodal completion. *Perception* 38, 17–29
9. Risko, E.F. *et al.* (2006) The ties that keep us bound: top-down influences on the persistence of shape-from-motion. *Conscious. Cogn.* 15, 475–483
10. Moore, C. and Cavanagh, P. (1998) Recovery of 3D volume from 2-tone images of novel objects. *Cognition* 67, 45–71
11. Bulthoff, I. *et al.* (1998) Top-down influences on stereoscopic depth-perception. *Nat. Neurosci.* 1, 254–257
12. Teufel, C. *et al.* (2018) Prior object-knowledge sharpens properties of early visual feature-detectors. *Sci. Rep.* 8, 1–12
13. Samaha, J. *et al.* (2018) Effects of meaningfulness on perception: alpha-band oscillations carry perceptual expectations and influence early visual responses. *Sci. Rep.* 8, 6606
14. Yuille, A. and Kersten, D. (2006) Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* 10, 301–308
15. Dehaene-Lambertz, G. *et al.* (2000) Electrophysiological correlates of phonological processing: a cross-linguistic study. *J. Cogn. Neurosci.* 12, 635–647
16. Kazanina, N. *et al.* (2006) The influence of meaning on the perception of speech sounds. *Proc. Natl. Acad. Sci.* 103, 11381–11386
17. Samuel, A.G. (1997) Lexical activation produces potent phonemic percepts. *Cogn. Psychol.* 32, 97–127
18. McClelland, J. *et al.* (2006) Are there interactive processes in speech perception? *Trends Cogn. Sci.* 10, 363–369
19. Cohen, L. *et al.* (2002) Language-specific tuning of visual cortex? Functional properties of the visual word form area. *Brain* 125, 1054–1069
20. Clark, A. (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav. Brain Sci.* 36, 181–204

21. Lupyan, G. (2015) Cognitive penetrability of perception in the age of prediction: predictive systems are penetrable systems. *Rev. Philos. Psychol.* 6, 547–569
22. O'Callaghan, C. *et al.* (2016) Predictions penetrate perception: converging insights from brain, behaviour and disorder. *Conscious. Cogn.* 47, 63–74
23. Harnad, S. (2005) Cognition Is Categorization. In *Handbook of Categorization in Cognitive Science* (Cohen, H. and Lefebvre, C., eds), pp. 20–45, Elsevier Science, San Diego, CA
24. Hsieh, P.-J. *et al.* (2010) Recognition alters the spatial pattern of FMRI activation in early retinotopic cortex. *Journal of Neurophysiology* 103, 1501–1507
25. Vandenbroucke, A.R.E. *et al.* (2016) Prior knowledge about objects determines neural color representation in human visual cortex. *Cereb. Cortex* 26, 1401–1408 (New York, N.Y.: 1991)
26. Pylyshyn, Z. (1999) Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behav. Brain Sci.* 22, 341–365
27. Dils, A.T. and Boroditsky, L. (2010) Processing unrelated language can change what you see. *Psychon. Bull. Rev.* 17, 882–888
28. Firestone, C. and Scholl, B. (2016) Cognition does not affect perception: evaluating the evidence for 'top-down' effects. *Behav. Brain Sci.* 39, 1–77
29. Luck, S.J. (2014) Overview of common ERP components. In *An Introduction to the Event-Related Potential Technique* (2 edn), pp. 71–118, MIT Press, Cambridge, Massachusetts
30. Rossion, B. and Jacques, C. (2011) The N170: understanding the time course of face perception in the human brain. In *The Oxford Handbook of Event-Related Potential Components* (Kappenman, E. and Luck, S., eds), pp. 115–142, Oxford University Press
31. Tanaka, J. and Curran, T. (2001) A neural basis for expert object recognition. *Psychological Science* 12, 43–47
32. Kutas, M. and Federmeier, K.D. (2011) Thirty years and counting: finding meaning in the N400 component of the event-related brain potential (ERP). *Annu. Rev. Psychol.* 62, 621–647
33. Rabovsky, M. *et al.* (2018) Modelling the N400 brain potential as change in a probabilistic representation of meaning. *Nat. Hum. Behav.* 2, 693–705
34. Abdel Rahman, R. and Sommer, W. (2008) Seeing what we know and understand: how knowledge shapes perception. *Psychon. Bull. Rev.* 15, 1055–1063
35. Rabovsky, M. *et al.* (2012) Depth of conceptual knowledge modulates visual processes during word reading. *J. Cogn. Neurosci.* 24, 990–1005
36. Boutonnet, B. and Lupyan, G. (2015) Words jump-start vision: a label advantage in object recognition. *J. Neurosci.* 32, 9329–9335
37. Feldman, N.H. *et al.* (2009) The influence of categories on perception: explaining the perceptual magnet effect as optimal statistical inference. *Psychol. Rev.* 116, 752–782
38. Forder, L. and Lupyan, G. (2019) Hearing words changes color perception: facilitation of color discrimination by verbal and visual cues. *J. Exp. Psychol. Gen.* 148, 1105–1123
39. Athanasopoulos, P. *et al.* (2009) The Whorfian mind: electrophysiological evidence that language shapes perception. *Commun. Integr. Biol.* 2, 332–334
40. Cibelli, E. *et al.* (2016) The Sapir-Whorf hypothesis and probabilistic inference: evidence from the domain of color. *PLoS One* 11, e0158725
41. Regier, T. and Kay, P. (2009) Language, thought, and color: Whorf was half right. *Trends Cogn. Sci.* 13, 439–446
42. Roberson, D. (2005) Color categories are culturally diverse in cognition as well as in language. *Cross-Cult. Res.* 39, 56–71
43. Anderson, W. *et al.*, eds (2014) *Colour Studies: A Broad Spectrum*, John Benjamins Publishing Company
44. Ocelák, R. (2016) "Categorical perception" and linguistic categorization of color. *Rev. Philos. Psychol.* 7, 55–70
45. Athanasopoulos, P. and Casaponsa, A. (2020) The Whorfian brain: neuroscientific approaches to linguistic relativity. *Cogn. Neuropsychol.* 0, 1–20
46. Winawer, J. *et al.* (2007) Russian blues reveal effects of language on color discrimination. *Proc. Natl. Acad. Sci. U. S. A.* 104, 7780–7785
47. Roberson, D. *et al.* (2008) Categorical perception of colour in the left and right visual field is verbally mediated: evidence from Korean. *Cognition* 107, 752–762
48. Maier, M. and Abdel Rahman, R. (2018) Native language promotes access to visual consciousness. *Psychol. Sci.* 29, 1757–1772
49. Thierry, G. *et al.* (2009) Unconscious effects of language-specific terminology on preattentive color perception. *Proc. Natl. Acad. Sci.* 106, 4567–4570
50. Forder, L. *et al.* (2017) Colour categories are reflected in sensory stages of colour perception when stimulus issues are resolved. *PLoS One* 12, e0178097
51. Boutonnet, B. *et al.* (2013) Seeing objects through the language glass. *J. Cogn. Neurosci.* 25, 1702–1710
52. Ozgen, E. and Davies, I. (2002) Acquisition of categorical color perception: a perceptual learning approach to the linguistic relativity hypothesis. *J. Exp. Psychol. Gen.* 131, 477–493
53. Zhou, K. *et al.* (2010) Newly trained lexical categories produce lateralized categorical perception of color. *Proc. Natl. Acad. Sci.* 107, 9974–9978
54. Grandison, A. *et al.* (2016) Chromatic perceptual learning but no category effects without linguistic input. *Front. Psychol.* 7, 731
55. Gilbert, A. *et al.* (2006) Whorf Hypothesis Is Supported in the Right Visual Field but Not the Left. *Proc. Natl. Acad. Sci. U. S. A.* 103, 489–494
56. Pilling, M. *et al.* (2003) Is color "categorical perception" really perceptual? *Mem. Cogn.* 31, 538–551
57. Roberson, D. *et al.* (2007) Categorical perception of facial expressions: evidence for a "category adjustment" model. *Mem. Cogn.* 35, 1814–1829
58. Drivonikou, G. *et al.* (2007) Further evidence that Whorfian effects are stronger in the right visual field than the left. *Proc. Natl. Acad. Sci.* 104, 1097–1102
59. Zhong, W. *et al.* (2018) Is the lateralized categorical perception of color a situational effect of language on color perception? *Cogn. Sci.* 42, 350–364
60. Mo, L. *et al.* (2011) Electrophysiological evidence for the left-lateralized effect of language on preattentive categorical perception of color. *Proc. Natl. Acad. Sci.* 108, 14026–14030
61. Ting Siok, W. *et al.* (2009) Language regions of brain are operative in color perception. *Proc. Natl. Acad. Sci.* 106, 8140–8145
62. Tan, L.H. *et al.* (2008) Language affects patterns of brain activation associated with perceptual decision. *Proc. Natl. Acad. Sci.* 105, 4004–4009
63. Kosslyn, S.M. *et al.* (1989) Evidence for two types of spatial representations: hemispheric specialization for categorical and coordinate relations. *J. Exp. Psychol. Hum. Percept. Perform.* 15, 723–735
64. Holmes, K.J. and Wolff, P. (2012) Does categorical perception in the left hemisphere depend on language? *J. Exp. Psychol. Gen.* 141, 439–443
65. Witzel, C. and Gegenfurtner, K.R. (2011) Is there a lateralized category effect for color? *J. Vis.* 11, 16
66. Brederoo, S.G. *et al.* (2019) Reproducibility of visual-field asymmetries: nine replication studies investigating lateralization of visual information processing. *Cortex J. Devoted Study Nerv. Syst. Behav.* 111, 100–126
67. Wright, O. *et al.* (2015) Whorfian effects on colour memory are not reliable. *Q. J. Exp. Psychol.* 68, 745–758
68. Souza, A.S. and Skóra, Z. (2017) The interplay of language and visual perception in working memory. *Cognition* 166, 277–297
69. Lupyan, G. and Spivey, M. (2008) Perceptual processing is facilitated by ascribing meaning to novel stimuli. *Curr. Biol.* 18, R410–R412
70. Lupyan, G. (2008) The conceptual grouping effect: categories matter (and named categories matter more). *Cognition* 108, 566–577
71. Gauthier, I. *et al.* (2003) The influence of conceptual knowledge on visual discrimination. *Cogn. Neuropsychol.* 20, 507–523
72. Edmiston, P. and Lupyan, G. (2015) What makes words special? Words as unmotivated cues. *Cognition* 143, 93–100
73. Lupyan, G. and Spivey, M. (2010) Making the invisible visible: Auditory cues facilitate visual object detection. *PLoS One* 5, e11452

74. Kouider, S. and Dehaene, S. (2007) Levels of processing during non-conscious perception: a critical review of visual masking. *Philos. Trans. R. Soc., B Biol. Sci.* 362, 857–875
75. Lupyan, G. and Ward, E.J. (2013) Language can boost otherwise unseen objects into visual awareness. *Proc. Natl. Acad. Sci.* 110, 14196–14201
76. Pasley, B.N. *et al.* (2004) Subcortical discrimination of unperceived objects during binocular rivalry. *Neuron* 42, 163–172
77. Noorman, S. *et al.* (2018) Words affect visual perception by activating object shape representations. *Sci. Rep.* 8, 14156
78. Raymond, J.E. *et al.* (1992) Temporary suppression of visual processing in an RSVP task: an attentional blink? *J. Exp. Psychol. Hum. Percept. Perform.* 18, 849–860
79. Chun, M.M. and Potter, M.C. (1995) A two-stage model for multiple target detection in rapid serial visual presentation. *J. Exp. Psychol. Hum. Percept. Perform.* 21, 109–127
80. Weller, P.D. *et al.* (2019) Semantic knowledge enhances conscious awareness of visual objects. *J. Cogn. Neurosci.* 31, 1216–1226
81. Foxe, J. and Simpson, G. (2002) Flow of activation from V1 to frontal cortex in humans - a framework for defining "early" visual processing. *Exp. Brain Res.* 142, 139–150
82. De Lange, F. *et al.* (2018) How do expectations shape perception? *Trends Cogn. Sci.* 22, 764–779
83. Clark, A. (2016) *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*, Oxford University Press
84. Hohwy, J. (2013) *The Predictive Mind*, Oxford University Press
85. Lupyan, G. and Thompson-Schill, S. (2012) The evocative power of words: activation of concepts by verbal and nonverbal means. *J. Exp. Psychol. Gen.* 141, 170–186
86. Chen, Y.-C. and Spence, C. (2011) Crossmodal semantic priming by naturalistic sounds and spoken words enhances visual sensitivity. *J. Exp. Psychol. Hum. Percept. Perform.* 37, 1554–1568
87. Yu, C.-P. *et al.* (2016) Searching for category-consistent features: a computational approach to understanding visual category representation. *Psychol. Sci.* 27, 870–884
88. Nako, R. *et al.* (2014) Rapid guidance of visual search by object categories. *J. Exp. Psychol. Hum. Percept. Perform.* 40, 50–60
89. Çukur, T. *et al.* (2013) Attention during natural vision warps semantic representation across the human brain. *Nat. Neurosci.* 16, 763–770
90. Brouwer, G.J. and Heeger, D.J. (2013) Categorical clustering of the neural representation of color. *J. Neurosci.* 33, 15454–15465
91. Hout, M.C. and Goldinger, S.D. (2015) Target templates: the precision of mental representations affects attentional guidance and decision-making in visual search. *Atten. Percept. Psychophys.* 77, 128–149
92. Roberson, D. *et al.* (2009) Thresholds for color discrimination in English and Korean speakers. *Cognition* 112, 482–487
93. Witzel, C. and Gegenfurtner, K.R. (2013) Categorical sensitivity to color differences. *J. Vis.* 13, 1
94. Webster, M.A. and Kay, P. (2012) Color categories and color appearance. *Cognition* 122, 375–392
95. Witzel, C. and Gegenfurtner, K.R. (2015) Categorical facilitation with equally discriminable colors. *J. Vis.* 15, 22
96. Gleitman, L. and Papafragou, A. (2012) New perspectives on language and thought. In *The Oxford Handbook of Thinking and Reasoning* (Holyoak, K. and Morrison, R., eds), pp. 543–568, Oxford University Press
97. Winawer, J. and Witthoft, N. (2020) Effects of color terms on color perception and cognition. In *Encyclopedia of Color Science and Technology* (Shamey, R., ed.), pp. 1–9, Springer, Berlin, Heidelberg
98. Huttenlocher, J. *et al.* (1991) Categories and particulars - prototype effects in estimating spatial location. *Psychol. Rev.* 98, 352–376
99. Bae, G.-Y. *et al.* (2015) Why some colors appear more memorable than others: a model combining categories and particulars in color working memory. *J. Exp. Psychol. Gen.* 144, 744–763
100. Ahissar, M. and Hochstein, S. (2004) The reverse hierarchy theory of visual perceptual learning. *Trends Cogn. Sci.* 8, 457–464
101. Sohoglu, E. and Davis, M.H. (2016) Perceptual learning of degraded speech by minimizing prediction error. *Proc. Natl. Acad. Sci.* 113, E1747–E1756
102. Cope, T.E. *et al.* (2017) Evidence for causal top-down frontal contributions to predictive processes in speech perception. *Nat. Commun.* 8, 2154
103. McWhorter, J.H. (2014) *The Language Hoax: Why the World Looks the Same in Any Language*, Oxford University Press
104. Wallisch, P. (2017) Illumination assumptions account for individual differences in the perceptual interpretation of a profoundly ambiguous stimulus in the color domain: "the dress". *J. Vis.* 17, 5
105. Firestone, C. and Scholl, B.J. (2015) Can you experience 'top-down' effects on perception?: the case of race categories and perceived lightness. *Psychon. Bull. Rev.* 22, 694–700
106. Adelson, E. (1993) Perceptual organization and the judgment of brightness. *Science* 262, 2042–2044
107. Perry, L.K. and Lupyan, G. (2013) What the online manipulation of linguistic activity can tell us about language and thought. *Front. Behav. Neurosci.* 7, 122
108. Lupyan, G. (2017) Changing what you see by changing what you know: the role of attention. *Front. Psychol.* 8, 553
109. Fodor, J. (1983) *The Modularity of Mind*, MIT Press
110. Brown, C. *et al.* (2008) Modulation of pain ratings by expectation and uncertainty: Behavioral characteristics and anticipatory neural correlates. *Pain* 135, 240–250
111. Pollo, A. *et al.* (2003) Placebo analgesia and the heart. *Pain* 102, 125–133
112. Medoff, Z.M. and Colloca, L. (2015) Placebo analgesia: understanding the mechanisms. *Pain Manag.* 5, 89–96
113. Colloca, L. *et al.* (2013) Placebo analgesia: psychological and neurobiological mechanisms. *Pain* 154, 511–514
114. Geuter, S. *et al.* (2017) Functional dissociation of stimulus intensity encoding and predictive coding of pain in the insula. *eLife* 6, e24770
115. Büchel, C. *et al.* (2014) Placebo analgesia: a predictive coding perspective. *Neuron* 81, 1223–1239
116. Wager, T.D. *et al.* (2004) Placebo-induced changes in fMRI in the anticipation and experience of pain. *Science* 303, 1162–1167
117. Watson, A. *et al.* (2006) Categories of placebo response in the absence of site-specific expectation of analgesia. *Pain* 126, 115–122
118. Benedetti, F. *et al.* (2003) Conscious expectation and unconscious conditioning in analgesic, motor, and hormonal placebo/nocebo responses. *J. Neurosci.* 23, 4315–4323
119. Colloca, L. and Miller, F.G. (2011) How placebo responses are formed: a learning perspective. *Philos. Trans. R. Soc., B Biol. Sci.* 366, 1859–1869
120. Lupyan, G. and Clark, A. (2015) Words and the world: predictive coding and the language-perception-cognition interface. *Curr. Dir. Psychol. Sci.* 24, 279–284
121. Wierzbicka, A. (1996) *Semantics: Primes and Universals*, Oxford University Press
122. Evans, N. and Levinson, S. (2009) The myth of language universals: language diversity and its importance for cognitive science. *Behav. Brain Sci.* 32, 429
123. Pullum, G.K. (1989) The great Eskimo vocabulary hoax. *Nat. Lang. Linguist. Theory* 7, 275–281
124. Majid, A. *et al.* (2018) Differential coding of perception in the world's languages. *Proc. Natl. Acad. Sci.* 115, 11369–11376
125. Zaslavsky, N. *et al.* (2018) Efficient compression in color naming and its evolution. *Proc. Natl. Acad. Sci.* 115, 7937–7942
126. Gibson, E. *et al.* (2017) Color naming across languages reflects color use. *Proc. Natl. Acad. Sci.* 114, 10785–10790
127. Rossman, I. and Goss, A. (1951) The acquired distinctiveness of cues: the role of discriminative verbal responses in facilitating the acquisition of discriminative motor responses. *J. Exp. Psychol.* 42, 173–182
128. de Saussure, F. (1916) *Course in General Linguistics*, Columbia University Press
129. Boleda, G. (2020) Distributional semantics and linguistic theory. *Annu. Rev. Linguist.* 6, 231–234
130. Elman, J. (2004) An alternative view of the mental lexicon. *Trends Cogn. Sci.* 8, 301–306
131. Baroni, M. *et al.* (2014) Don't count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of the 52nd Annual Meeting of the*

Association for Computational Linguistics Baltimore, MD, pp. 238–247
132. Dimitrova-Radojichikj, D. (2015) Concepts of colors in children with congenital blindness. *J. Spec. Educ. Rehab.* 16, 7–16
133. Shepard, R.N. and Cooper, L.A. (1992) Representation of colors in the blind, color-blind, and normally sighted. *Psychol. Sci.* 3, 97–104
134. Lenci, A. *et al.* (2013) BLIND: a set of semantic feature norms from the congenitally blind. *Behav. Res. Methods* 45, 1218–1233
135. Marmor, G.S. (1978) Age at onset of blindness and the development of the semantics of color names. *J. Exp. Child Psychol.* 25, 267–278
136. Kim, J.S. *et al.* (2019) Knowledge of animal appearance among sighted and blind adults. *Proc. Natl. Acad. Sci.* 116, 11213–11222
137. Lewis, M. *et al.* (2019) Distributional semantics as a source of visual knowledge. *Proc. Natl. Acad. Sci.* 116, 19237–19238
138. Lupyan, G. and Lewis, M. (2017) From words-as-mappings to words-as-cues: the role of language in semantic knowledge. *Lang. Cogn. Neurosci.* 34, 1319–1337
139. Roebuck, H. and Lupyan, G. (2020) The Internal Representations Questionnaire: measuring modes of thinking. *Behav. Res. Methods* https://doi.org/10.3758/s13428-020-01354-y
140. Cichy, R.M. *et al.* (2012) Imagery and perception share cortical representations of content and location. *Cereb. Cortex* 22, 372–380
141. Dijkstra, N. *et al.* (2017) Vividness of visual imagery depends on the neural overlap with perception in visual areas. *J. Neurosci.* 37, 1367–1373
142. Kosslyn, S.M. (2005) Mental images and the brain. *Cogn. Neuropsychol.* 22, 333–347
143. Horikawa, T. and Kamitani, Y. (2017) Generic decoding of seen and imagined objects using hierarchical visual features. *Nat. Commun.* 8, 1–15
144. Dils, A.T. and Boroditsky, L. (2010) Visual motion aftereffect from understanding motion language. *Proc. Natl. Acad. Sci.* 107, 16396–16400
145. Maier, M. *et al.* (2020) Time course and shared neurocognitive mechanisms of mental imagery and visual perception. *bioRxiv.* https://doi.org/10.1101/2020.01.14.905885 2020.01.14.905885
146. Tseng, C. *et al.* (2016) A computational investigation of the Sapir-Whorf Hypothesis: the case of spatial relations. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (Papafragou, A., ed.), pp. 2231–2236, Cognitive Science Society, Austin, TX