DOI: 10.1111/mila.12159

WILEY

ORIGINAL ARTICLE

Perceptual expansion under cognitive guidance: Lessons from language processing

Endre Begby

Department of Philosophy, Simon Fraser University, Burnaby, BC, Canada

Correspondence
Endre Begby, Department of Philosophy, Simon Fraser University, 8888 University Drive, Burnaby, BC V5A 1S6, Canada.
Email: endre.begby@gmail.com

This paper aims to provide an empirically informed sketch of how our perceptual capacities can interact with cognitive processes to give rise to new perceptual attributives. In section 1, I present ongoing debates about the reach of perception and direct focus toward arguments offered in recent work by Tyler Burge and Ned Block. In section 2, I draw on empirical evidence relating to language processing to argue against the claim that we have no acquired, culture-specific, high-level perceptual attributives. In section 3, I turn to the cognitive dimension; I outline how cognitive procedures (including conceptual representation and explicit inference) can be involved in the acquisition of what ought to, nonetheless, be recognized as genuinely perceptual capacities. Finally, in section 4, I argue for the importance of distinguishing these conclusions from more familiar and radical claims about rampant "cognitive penetration" into the perceptual domain.

KEYWORDS
cognitive penetration, high-level perception, perception, perception–cognition interface, perceptual constancies, perceptual learning

1 | INTRODUCTION

A central dividing line in the philosophy of perception concerns the demarcation of perceptual capacities from conceptual thought and other higher cognitive processes. On this question, there is, one would think, a whole spectrum of viable positions. Nonetheless, it seems that most contributions tend to cluster near one pole or the other. Thus, on the one hand, there are those who favor a more *austere* notion of perception, according to which our perceptual faculties are sharply demarcated from higher cognitive functions and serve to output representations of only a relatively narrow range of low-level properties, such as (in the case of vision) edges, textures, and, perhaps, colors (c.f., Dretske, 1995; Tye, 1995; Prinz, 2005; Brogaard, 2013; Reiland, 2014). On the other hand, there are those who favor a more *liberal* notion of perception, according to which perceptual systems can interact in intricate ways with higher cognitive functions and are, as a result, capable of outputting representations of much more complex properties, perhaps including the properties that distinguish everyday objects, such as toasters and teacups.[1] Both sides will claim distinctive sorts of evidence in support of their views: typically, austerists will favor evidence from contemporary perceptual psychology, whereas liberals will appeal to phenomenological data (for a prominent example, consider the "method of phenomenal contrast" prescribed in Siegel, 2006 and Siegel, 2010).

An important recent contribution to these debates is Tyler Burge's *Origins of Objectivity* (Burge, 2010). Burge's account of perception is shaped around three general ideas familiar from contemporary perceptual psychology (c.f., Fodor, 1983; Pylyshyn, 1984, 1999): (a) *Domain specificity*: each perceptual system (for instance, vision) delivers representations of only a relatively small number of low-level attributes (for instance, shape, brightness, color); (b) *Encapsulation*: perceptual systems function relatively independently of input from each other and from higher-level cognitive faculties; and (c) Perceptual systems are *shared across species*: for instance, vision functions in the same way, by and large, across most mammalian species (Burge, 2010, pp. 101–102).[2]

In Burge's terminology, the question of what range of properties can be represented in perception is settled in terms of the question of what range of "perceptual attributives" we have at our disposal.[3] This is an empirical question, but Burge is adamant that the empirical evidence dictates a relatively minimal and austere answer. It is true that Burge does not deny the existence of high(er)-level perceptual attributives altogether. For instance, the attributive *integrated body* receives lengthy treatment in his book (Burge, 2010, pp. 437–470), and in response to a recent paper by Ned Block (Block, 2014), he concedes that certain aspects of face perception, as well as some notions relating to causality and agency, may also be perceptual in character.[4] But even with these amendments, it seems fair to say that Burge's view of the reach of perception is decidedly on the austere end of the spectrum. Specifically, he holds firmly that we do not have perceptual attributives for the properties that distinguish everyday objects, such as "baseball bats, CD-players, hybrid autos" (Burge, 2010, p. 101). Thus, although we certainly have perceptual *beliefs* about these objects, we do not really perceive them as such. Rather, what we perceive, in the strict and proper sense, are simply their low-level properties, such as colors, shapes, and textures. Representation of these more complex objects and properties "depends on capacities that go beyond the perceptual system proper" (Burge, 2010, p. 101) and is accordingly better thought of in terms of "hybrids" of perception and conceptual representation (e.g., Block, 2014, p. 566; Burge, 2010, p. 546). Further, and this will be my entry point in this paper, this approach also appears to cast serious aspersions on the idea that we

[1] Cf. Siegel, 2006 and Siegel, 2010; Bayne, 2009; Fish, 2013. A number of important papers on this topic are collected in Hawley & Macpherson, 2011 and Zeimbekis & Raftopoulos, 2015. (I should add that my argument bears on disputes about the reach of perceptual *content*—that is, what range of properties receives representation in perceptual experience—and not merely about the reach of perceptual phenomenology as such [c.f., Masrour, 2011]. Many thanks to an anonymous referee for raising this issue.)

[2] For a detailed overview, see Begby, 2011. It is worth noting that Burge emphatically does *not* advocate such a view because he believes it constitutes a useful demarcation for theoretical and practical purposes. Instead, Burge's motivation is metaphysical, and he is adamant that his demarcation traces the boundaries of a genuine natural kind, a "cut in nature" (Burge, 2010, p. 9, p. xii).

[3] Briefly, perceptual attributives constitute the general (as opposed to singular) elements of perception, whose role is "to group or categorize particulars by attributing some indicated kind, property, or relation to them" (Burge, 2010, p. 83). More fully, "[a] perceptual attributive is an aspect of perceptual representational content that functions to indicate a repeatable type and to group or characterize purported particulars as being of that type" (Burge, 2010, p. 380). If we will, we can think of them as having a role in the nonpropositional structure of perceptual representation corresponding to that of concepts in propositional thoughts.

[4] "[T]hese conjectures," he nonetheless warns, "are not empirically secure" (Burge, 2014, p. 583).

might be in possession of any *culture-specific*, high-level perceptual attributives of the sort that would be required for us to perceive—in the strict and proper sense—objects such as CD players, baseball bats, and hybrid autos (see, e.g., Burge, 2010, p. 101; Burge, 2014, p. 575; Block, 2014, p. 560).[5]

I mentioned above how different approaches to these issues typically appeal to different bodies of evidence. Interestingly, Burge and Block both seem to hold that *only* appeals to phenomenology could motivate the idea that we might be in possession of such culture-specific perceptual attributives. Burge finds these appeals to be "deeply wrong-headed" and "epistemically worthless." He writes: "Only sophisticated use of experimental evidence bears on these issues in a way that goes beyond uninformed playing" (Burge, 2014, p. 583). Block is a shade less dismissive: he seems to allow that conclusions drawn from phenomenological data may have a certain kind of *prima facie* validity. Nonetheless, he believes it can be shown—on an empirical basis—that the properties they invoke are "recognitionally coextensive" with clusters of lower-level perceptual properties (Block, 2014, pp. 562–563), thereby allowing us to maintain, in line with Burge's conclusion above, that what we really perceive—strictly speaking—are just the familiar low-level attributes, such as "color, shape, and texture" (Block, 2014, p. 560).

In this paper, I aim to develop a case for a (more) liberal account of perception by drawing precisely on empirical rather than phenomenological considerations. I will not go so far as to argue that we can perceive toasters and teacups as such (though I would not rule out that we can). Like Burge and Block, I am primarily interested in determining in which direction the available empirical evidence is pointing. Accordingly, my aim is to provide an empirically informed sketch of how our perceptual capacities can interact with cognitive processes to give rise to new culture-specific perceptual attributives in ways that would be quite at odds with the conclusions of Burge and Block.[6] In section 2 of this paper, I draw on empirical evidence relating to language processing to argue against the claim that we have no culture-specific perceptual attributives. In section 3, I turn to the cognitive dimension; I outline how cognitive procedures (including conceptual representation and explicit inference) can very well be involved in the acquisition of what ought, nonetheless, to be recognized as genuinely perceptual capacities (fluent reading is a case in point). Finally, in section 4, I argue for the importance of distinguishing these conclusions from more familiar and radical arguments for rampant "cognitive penetration" into the perceptual domain. As I point out, my conclusions are actually consistent with the assumption, central to Burge and Block, that the perceptual domain is *synchronically* well-demarcated with respect to higher cognitive processes: *at any given time*, the perceptual system operates only with such attributives as it has in stock. Rather, my point is that there are processes straddling the perception–cognition boundary that can add to the stock of perceptual attributives *over time.*

2 | EXAMPLES FROM LANGUAGE PROCESSING

Our guiding question, then, is whether there is evidence that we are in possession of culture-specific, high-level perceptual attributives. To shed light on this question, I propose that we turn to our

[5] See Begby, 2011, for an earlier attempt to raise concerns about this sort of view, to which Block, 2014, is in part a response.

[6] While I think my argumentative strategy is unique in the contemporary literature, it is not without historical precedent. Rather, it is importantly inspired by Thomas Reid's distinction between original and acquired perception and, in particular, his claim that by deploying innate, low-level perceptual capacities in concert with cognitive capacities, we can come to extend our perceptual horizons in important ways. On this, see Reid, 1764, VI: 20–23; Reid, 1785, II: 21–22.

perceptual facility in dealing with language, spoken and written. As such, a prefatory note is in order: language perception is not central to Burge's book and is mentioned in passing only a few times. Here is one instance:

> In humans and higher animals, perception interfaces with conception and belief in complex ways. Nevertheless, the processes of perceptual systems, even in humans, are *relatively* independent of higher-level cognitive states. (Language perception is a special case and requires further qualification.) (Burge, 2010, p. 102)

Note that Burge says nothing further to specify in what sense language perception is a "special case" or to cast light on the nature and extent of the "further qualification" required. But given the context, it may be plausible to assume that he has in mind well-known effects, such as the phonemic restoration effect and the McGurk effect. The former offers a striking instance of so-called top–down influence on perceptual processing: subjects hear a speech signal as complete even though parts of the signal have been removed and replaced with noise; the perceptual system appears to "fill in" the missing phonemes on the basis of projections of semantic cohesion and completeness. By contrast, the McGurk effect illustrates cross-modal interference: subjects are played two video clips in sequence, the first showing a person voicing a particular syllable (say, /ba/) with matched audio. In the second clip, the audio stays the same (/ba/), but the video is changed to someone uttering a different syllable (say, /ga/). The visual perception of lips forming the syllable /ga/ appears to override the actual auditory input, leading the subject to perceptually represent a voicing of yet a third syllable, such as /da/. Notably, this effect persists even after subjects are informed of its existence.

These examples, interesting and systematic as they appear to be, may nevertheless be compatible with Burge's sweeping and programmatic claims that perceptual processing is "relatively" encapsulated and operates in "relative" independence of higher-order cognitive states. I will not pursue the matter here.

2.1 | Phoneme perception

However, language processing also provides examples that offer a more sweeping challenge to austere views of the reach of perception. As has been known for some time (see, e.g., Liberman, Cooper, Shankweiler, & Studdert-Kennedy, 1967; Fodor, Bever, & Garrett, 1974), there is a deep mismatch between the acoustic signal and our perceptual representation of speech. For instance, we perceive speech as segmented into discrete units such as syllables and words. As is revealed by spectrogram analysis, however, this segmentation fails to correlate with any properties found in the acoustic signal itself. Word boundaries are largely nonexistent, and many phonemic elements that we perceive as linearly ordered are in fact coarticulated. Finally, and most importantly for the argument that follows, speech sounds are highly variant, not just with respect to who is speaking—gender, age, dialect, and so forth—but also with respect to the surrounding phonemic context; for instance, the physical realization of /d/ depends on what sort of vowel it is followed by. Yet, we perceive the /d/ in the same way in words such as "dib," "dab," and "dub." Conversely, other speech sounds may be perceptually represented as distinct, again depending on phonemic context, even though their physical realization is very nearly identical.

I believe it is eminently clear that speech processing involves distinctive perceptual capacities at a number of levels. More specifically, I believe it is clear that phonemes should count as perceptual attributives in Burge's sense: pertinently, they involve significant constancy mechanisms, grouping together distinct physical stimuli as instances of a single kind.[7]

A number of points suggest themselves right away. First, phonemes are also, at least in some sense, *high-level* perceptual attributives. Neither Burge nor Block offers any strict definition of what is meant by "high-level perceptual attributive." But if *integrated body* is high level, as Burge and Block argue, then presumably they have in mind a contrast with the low-level inventory of particular perceptual modalities. Block suggests, for vision, a low-level inventory of "shape, spatial relations (including position and size), geometrical motion, texture, brightness, and color" (Block, 2014, p. 560). Presumably, a comparable inventory for audition will center on features such as pitch, timbre, and loudness. Phonemes are not on that list, nor do they correspond to clusters of such low-level attributives. Arguably, phonemes are not even to be identified with complex objects, such as sounds. Instead, they might be better thought of in terms of *equivalence classes of sounds* (c.f., O'Callaghan, 2014). And that, I take it, is clearly a high-level attributive.

Second, phoneme perception is plausibly also a *species-specific capacity.*[8] (Other species may learn to discriminate a range of human speech sounds, but not by way of the sorts of perceptual mechanisms that we deploy [see, e.g., Trout, 2001; Pinker & Jackendoff, 2005, section 2.2].) This suggests the following point, which I will make further use of below: there remains, to be sure, a sense in which Burge might be right to say that hearing *does* function—by and large—in the same way across many species, certainly mammalian species. But in discussing the reach of perception, it is a mistake to look only at the major sensory pathways and their generic or archetypal representational outputs. We must also consider how they operate with respect to specific tasks. And, as the example of phoneme grouping shows, the perceptual capacity of hearing *as deployed toward the specific task of processing speech* has peculiarities unique to our species, peculiarities that ought to be called "perceptual" in their own right.

So far, our conclusion is quite modest, and I doubt that either Burge or Block would ever think to deny it. One key point here is the following: it may well be argued that phoneme grouping does not involve *acquired* perceptual capacities in the sense that perceptual liberals may be looking for. A familiar story from psycholinguistics has it that we are born with the capacity to discriminate all possible speech sounds deployable in human language and that we only come to "prune" these abilities with exposure to our first language (for an overview, see Werker & Tees, 2002). If this is correct, then phoneme perception does not support the claim that we can acquire new perceptual forms, however much it may be unique to our species.

My third point, then, is that there may, nonetheless, be a sense in which speech processing can be said to involve *culture-specific* perceptual attributives. Here is how: linguists define an "allophone" as the set of acoustically different sounds that speakers of particular languages perceptually group together as instantiating the same phoneme. For example, aspirated and nonaspirated /t/s are acoustically very different but are perceptually represented by monolingual English speakers as "the

[7] See O'Callaghan, 2011, for an excellent overview of these issues. See also Burge, 2010, pp. 233–234, 397, 399, 408 and 413 for the role of constancy mechanisms in demarcating the perceptual domain. We might note, in passing, that the literature on speech perception is sometimes couched in terms of the related but distinct phenomenon of "categorical perception." Categorical perception occurs when subjects display a significantly better ability to discriminate stimuli belonging to different categories (e.g., green, blue) than for discriminating stimuli belonging to the same category (e.g., different shades of blue) even when these stimuli are evenly spaced along a continuum. (See, e.g., Goldstone & Hendrickson, 2010, for an overview.) While it is plausible that there could be numerous category effects in perception that are *not* rooted in constancy mechanisms, this should in no way obscure the fact that speech perception also involves a number of important and distinctive constancy mechanisms, such as the ones that lead us to represent highly variable acoustic stimuli as instantiating the same phoneme. Unfortunately, however, the psycholinguistics literature is not always clear on this distinction (on this, see, e.g., Kluender & Kiefte, 2006, pp. 171–172).

[8] As Burge acknowledges, en passant, in an unrelated discussion (Burge, 2010, p. 320).

same sound," plausibly because the distinction is not semantically relevant in English. In Spanish, by contrast, the distinction between aspirated and nonaspirated /t/s does correspond to semantically relevant distinctions. Accordingly, they are perceptually represented by first-language Spanish speakers as distinct sounds (see, e.g., Jusczyk, 2003; Casserly & Pisoni, 2010).

To see the significance of this, consider the perhaps more familiar case of the failure of the /l/–/r/ contrast in Japanese: the ordinary learning situation for monolingual Japanese speakers does not lead them to produce perceptual representations that distinguish between these sounds. Even mature speakers who have attained a high degree of fluency in English, well past the point where they can confidently and correctly *articulate* the different sounds in the appropriate contexts, nonetheless continue to have difficulties *hearing* the difference in the speech of others (c.f., Goto, 1971; Sheldon & Strange, 1982). Thus, the English /l/–/r/ contrast might well be taken to illustrate a sense—though, admittedly, a thin sense—of a "culture-specific perceptual attributive," of the sort that Burge and Block cast aspersions on. Even as we all deploy universally shared innate capacities to come to terms with the perceptual challenges of human speech, exposure to different learning environments gives rise to predictably different repertoires of perceptual representation. The limitations of each repertoire are not easily overridden, even by voluminous subsequent exposure to a different learning environment.[9]

2.2 | Grapheme perception

For a clearer and more decisive case of culture-specific perceptual attributives, however, we can turn to the role of graphemes in reading. Defined on analogy with phonemes in spoken language, graphemes are the minimal units of written language that can make a semantic difference. Accordingly, each letter of the alphabet is a grapheme. But many languages will also have complex, multiletter graphemes, such as "ea" as it occurs in the English word "meat."

The question of what role, if any, graphemes play in the perceptual processes involved in reading has been with us at least since the work of Eleanor Gibson and associates in the early 1960s (see, e.g., Gibson, Pick, Osser, & Hammond, 1962). A key issue in current research on this topic is the problem of stimulus invariance, which the visual system must grapple with during reading.[10] It is worth noting that reading, in general, does not reproduce all of the perceptual challenges involved in speech processing: for instance, the stimulus is, for the most part, both relevantly segmented and linearly ordered. But it does require solving a problem of invariance quite similar to that involved in phoneme perception. That is, any letter can be physically realized in an indefinite number of ways, all of which appear to be perceptually processed by fluent readers as the "same grapheme" (Dehaene, 2010, pp. 18–20; McCandliss, Cohen, & Dehaene, 2003, p. 295). More specifically, while the low-level perceptual capacities involved in reading are highly sensitive to, say, font, case, tilt, and color, these features are, in most cases, informationally irrelevant to the computational task at hand. Reading, as a perceptual task, involves filtering out these features in the construction of a

---

[9] We may also note, in passing, that while all the classic studies focus on speech, it may be a mistake to conclude that these effects are the exclusive remit of auditory perception. In fact, there is emerging evidence that similar effects might also be operative in sign language perception. On this, see Emmorey, McCullough, & Brentari, 2003; Baker, Idsardi, Golinkoff, & Petitto, 2005; Best, Mathur, Miranda, & Lillo-Martin, 2010.

[10] Another approach focuses specifically on multiletter graphemes, claiming to show that subjects have a harder time detecting the presence of, for instance, the letter "A" as it occurs in the context of the word "BEACH" than they do when it occurs in the context of the word "PLACE" (see, e.g., Rey, Ziegler, & Jacobs, 2000; Marinus & de Jong, 2011). There is considerable uncertainty about just how robust these results are (c.f., Bolger, Borgwaldt, & Jakab, 2009; Lupker, Acha, Davis, & Perea, 2012). But even if the results should turn out to be replicable, this approach offers, at best, a more restricted and more circumstantial source of evidence for the perceptual role of graphemes than the approach I will outline below.

more abstract, higher-level representation—namely, the grapheme. Arguably, this desensitization reaches well beyond standard constancy mechanisms relating to shape and size and instead indicates an attunement of the visual system to "the specific demands of reading" (McCandliss et al., 2003, p. 295; see also Pegado et al., 2014).

To show just how deep this "attunement" runs, it will be helpful to look in more detail at one particularly telling manifestation of the invariance problem, namely, that of case. In a pioneering study, Polk and Farah (2002) conducted experiments with anomalous stimuli, such as "ElEpHaNt," demonstrating that these stimuli activate the brain's Visual Word Form Area (a part of the fusiform gyrus area, otherwise famous for its role in face recognition) in just the same way as regular, case-consistent stimuli do. It seems, then, that at some point in relatively early stages of visual word processing, the visual system has learned to filter out these otherwise very salient features of the visual stimulus in the construction of a more abstract perceptual attributive. To state the point more fully, it will be helpful to introduce the concept of an *allograph*—defined on analogy with "allophone" as per section 2.1 above. An allograph is the set of variant forms that a grapheme might take while still being processed by the visual system as "the same grapheme." The observation on offer in these experiments, then, is that the Visual Word Form Area treats, for instance, "a" and "A" as instantiating the same grapheme in much the same way as the auditory system of English speakers treats aspirated and nonaspirated /t/s as instantiating the same phoneme.

Further evidence along these lines comes from experiments conducted by Dehaene and associates. A well-known effect in neuropsychological studies is that of "repetition suppression," whereby repetitions of particular visual stimuli (e.g., "hotel"—"hotel") will elicit a distinctive pattern of decreased neuronal response. Strikingly, Dehaene et al. found that essentially identical patterns of repetition suppression occur in the Visual Word Form Area even when the stimulus word is repeated in a different case (e.g., "hotel"—"HOTEL") (Dehaene et al., 2001, p. 755; Dehaene, 2009, pp. 90–92; Hannagan, Amedi, Cohen, Dehaene-Lambertz, & Dehaene, 2015, p. 2.) This is strong evidence that even at early stages of visual processing, graphemes are represented as the "same" stimulus irrespective of case despite very salient differences in low-level perceptual features. As Dehaene observes (Dehaene, 2009, p. 161), "as neurons respond to increasingly complex combinations of visual features, their invariance is greater and they therefore drop other distinctions that are irrelevant for reading. The discrepancy between uppercase and lowercase letters, for instance, ceases to matter at an early stage in the visual stream."

Now, if this effect were restricted to particular case pairings (e.g., "O"/"o," "U"/"u"), we might seek to account for it in terms of simpler and more generic constancy mechanisms, such as size constancy. This would be highly congenial to Burge's approach. But that is precisely what we do not find: as Dehaene's experiments show, case pairings such as "G"/"g," "E"/"e," and "R"/"r" give rise to the same pattern of repetition suppression despite having little or no commonality in terms of low-level perceptual features.

Reading, in short, appears to be a very different perceptual exercise than visually tracing marks on the page, although, of course, there remains a sense in which we could not read were we not also able, concurrently, to perceive these marks. Accordingly, there seems to be a strong empirical case to be made that graphemes, and not simpler visual features such as lines, curves, and dots, are the proper objects of visual perception as deployed toward the task of reading, much as phonemes, rather than acoustic waveforms, are the proper object of auditory perception as deployed toward the task of processing speech.

These observations provide precisely what was still lacking in the argument from phoneme perception offered in section 2.1 above. First, I assume it will be granted that the perceptual capacities involved in reading are *acquired by learning* in a way the perceptual capacities involved in speech processing may not be.[11] I will have more to say about this learning process in section 3 below. Further, let us also note that the perceptual capacities involved in reading are *clearly* culture-specific: writing is a recent innovation, scarcely more than 5,000 years old. While literacy is spreading, an estimated one-fifth of the world's population remains unable to read or write. Only a handful of the world's 6,000 or so extant languages have a written standard in widespread use. Finally, while normal linguistic development requires no more than sustained exposure to spoken language, learning to read requires explicit instruction over a sustained period of time (for an overview, see Fernandez & Smith Cairns, 2010, pp. 93–94).

In summary, a closer look at our perceptual facility in dealing with language, spoken and written, reveals strong *empirical*, and not just phenomenological, evidence that we are in possession of culture-specific, high-level perceptual attributives. The observation that case discrepancies are filtered out even at relatively early stages of visual processing provides a perfect encapsulation of the point. For, as Dehaene points out, what ties together case pairs such as "g" and "G" is nothing but an arbitrary cultural convention. Accordingly, the fact "that neurons respond in the same way to the shapes 'g' and 'G' cannot be attributed to an innate organization of vision. It necessarily results from a learning process that has incorporated cultural practices into the appropriate brain networks" (Dehaene, 2009, p. 95).

3 | A ROLE FOR COGNITION?

Admittedly, however, phonemes and graphemes are still pretty low-level phenomena and will hardly whet the appetite of those who usually favor a more liberal account of the reach of perception. They are certainly not at the level that explicitly worries Burge and Block—"CD-players, baseball bats, hybrid autos." I have no argument that bears directly on the question of whether we have perceptual attributives of such objects. But I believe we can remove some of the obstacles that are typically taken to stand in the way of more liberal accounts by considering the involvement of higher-level *cognitive* capacities in the acquisition of reading.

Burge, as we have seen, dismisses more liberal views on the grounds that representation of objects such as CD players, baseball bats, and hybrid autos must depend on "capacities that go beyond the perceptual system proper." As stated, this is probably true. But it does not yet constitute an argument that we could not perceive such objects. The liberal account holds that our perceptual repertoires may be significantly augmented precisely as a result of interactions between our "perceptual capacities proper" and our cognitive capacities. On pain of begging the question, we could not simply stipulate that the involvement of cognitive capacities alone is evidence that the resulting representations are not perceptual. Here, as elsewhere, we would do well to follow where the empirical evidence leads.

In the recent exchange with Block, Burge looks to the role of attention and memory in diagnosing what goes wrong in liberal accounts of perception. He writes:

> One common error is to overlook the role of attention in grouping low-level attributes that perceptually trigger non-perceptual higher-level attributions. … [I]n the acquisition

[11] Indeed, as is emphasized in much of the recent literature on grapheme–color synesthesia (e.g., Watson, Blair, Kozik, Akins, & Enns, 2012), the learning process can be immensely difficult, especially in languages with highly irregular phoneme–grapheme correspondence, such as English.

> of expertise, conceptual higher-level attributives can cause attention to and new grouping of specific low-level attributes. (Burge, 2014, p. 582)

Further:

> [a] complimentary error is to overlook non-perceptual-attributive formation in long-term memory. […] Claims that perception is penetrated by higher-level attributives often just assume that higher-level attributives that are applied on the basis of perception are perceptual. No effort is made to rule out standard explanations that locate them in cognitive memory. (Burge, 2014, pp. 582–583)

This argument sketch may plausibly tell against many of the standard examples offered by perceptual liberals. How does it bear on the examples offered in section 2? Recall the point I made above, that in assessing the reach of perception, we cannot simply look at major sensory pathways (vision, audition, etc.) and their generic or archetypal outputs, but must also look at the way that these capacities are harnessed toward meeting particular specialized tasks, such as speech processing and reading. Thus, the perceptual capacity of audition *as deployed toward the task of processing speech* involves a distinctive set of perceptual attributives that cannot be reduced to the low-level inventory of auditory perception. Similarly, the perceptual capacity of vision *as deployed toward the task of reading* involves a distinctive set of perceptual attributives that cannot be reduced to the low-level inventory of visual perception.

Now, the perceptual capacities involved in processing speech may well be *sui generis.* Although I above described the process as *resulting* in culture-specific, high-level perceptual attributives, the process by which those attributives are generated involves no real role for higher cognitive faculties (at least in first language acquisition; second language acquisition may well be different in important respects). Instead, it is largely automatic, subpersonal, and innately programmed.

The case of reading, however, is importantly different and points toward a serious role for cognitive capacities—specifically, attention and memory—in giving rise to new high-level perceptual attributives. Learning to read certainly requires making novel use of low-level visuoperceptual capacities—for example, those that allow us to recognize lines on paper. But crucially, this is not all that is involved. Learning to read also requires harnessing other high-level perceptual capacities, in particular phonological capacities (see, e.g., Frith, 1985; Dehaene, 2009, pp. 199–204).[12] Finally, and most tellingly, since learning to read takes explicit instruction over a substantial period of time, it also centrally involves the deployment of higher *cognitive* functions, such as memory and attention. In early stages, learners support themselves on explicit inferential routines that map low-level visual shapes to abstract phonological forms. But at some point, these inferential routines are dropped: the subject is now a fluent reader. This transition marks the onset of what Uta Frith (1985) calls the "orthographic" stage, where written stimuli are no longer processed by way of explicit routines mapping graphemes to phonemes. At this point, the subject has acquired a new range of perceptual attributives, namely, the graphemes of his or her written language.

Contrary to Burge, then, capacities "beyond the perceptual system proper" may well be involved in the acquisition of new *perceptual* attributives. The fact that cognitive faculties, such as attention and memory, are involved in the acquisition of the capacities does not tell against the claim that the capacities acquired are perceptual capacities.

---

[12] Note, in particular, that deficits in phonological processing remain one of the most important predictors of developmental dyslexia. I will return to this point later.

To see this in more detail, consider Block's notion of "secondary seeing," which he proposes in an attempt to capture our phenomenological sense of perceiving complex high-level objects, such as New College. While primary seeing involves "the application of a visual attributive to a 'visual object', that is, an object that is itself picked out by a demonstrative element in a percept [….], [s]econdary seeing involves hybrids of visual attributives and concepts applied to objects of primary seeing and complexes of them in states that put together perception with perceptual judgment" (Block, 2014, p. 566).

Is fluent reading a case of "secondary seeing" in this sense? Certainly, we *do* have concepts of graphemes. But on its own, this is hardly grounds for concluding that we do not also have perceptual attributives of graphemes: after all, we may also have concepts of phonemes (not to mention concepts of colors, shapes, integrated bodies, etc.). If there is a difference, it lies in the normal order of acquisition: our concepts of phonemes, colors, etc. are typically acquired only after, and as a result of, full development of the relevant perceptual capacities. By contrast, in the development of reading, it is plausible that we acquire concepts of graphemes first (since explicit instruction is involved) and the perceptual abilities only subsequently. Nonetheless, this is fully explicable on the model that I am suggesting. These concepts figure in the inferential routines that support early stages of learning. But fluency is achieved only when these inferential routines are compressed and automatized in subpersonal, computational procedures that are directly triggered by particular kinds of sensory stimuli. At that point, perceptual learning has, via its interface with cognitive capacities, given rise to a new set of perceptual attributives. The mere fact that concepts are involved in this process of learning in no way tells against the conclusion that the outcome of the process is a significantly expanded perceptual repertoire.

Moreover, I do not think there is any empirically plausible case to be made that fluent reading consists in, say, a subject's framing a continuous succession of perceptual judgments relating clusters of low-level visual shapes—lines, curves, and dots—to abstract phonological forms (judgments, presumably, in which graphemes would figure as concepts, in accordance with Block's "hybrid" model). Interestingly, this issue has come up previously, in a debate between Paul Churchland and Jerry Fodor. Churchland (1988, p. 177) pointed precisely to reading as an instance of an acquired perceptual capacity that goes well beyond anything that our visual systems might have evolved to handle. Fodor's starkly dismissive reply essentially lumps this rather modest claim together with more extravagant liberal claims that we can visually perceive (as such) automobiles and other everyday artifacts:

> In recent centuries we have learned to [visually] perceive automobiles [...]. Now the eyes were not evolved for the instantaneous perception of those complex structures. So doesn't their acquired mastery illustrate the highly sophisticated and supernormal capacities that learning can produce in perception? Fiddlesticks. Churchland needs, and doesn't have, an argument that the visual perceptual capacities of people who can read (or, mutatis mutandis, people who can automobile spot) differ in any interesting way from the visual perceptual capacities of people who can't. In precisely what respects does he suppose illiterates to be visually incapacitated? The old story is: you read (spot automobiles) by making educated inferences from properties of things that your visual system was evolved to detect; shape, form, color, sequence and the like. Churchland offers no evidence that educating the inferences alters the perceptual apparatus. (Fodor, 1988, p. 194)

Accordingly, Fodor staunchly maintains that reading proceeds by inference from perception of familiar low-level visual properties, such as lines, curves, and dots. I submit that if they are to maintain their stance against acquired culture-specific, high-level perceptual attributives, Burge and Block will have to embrace a similar story. Like Fodor, they are certainly entitled to ask for evidence. But that evidence is now within our reach: there is strong, converging evidence from developmental psychology and neuropsychology that graphemes play a critical role in fluent reading. But graphemes are not reducible to low-level visual properties, such as shape, any more than phonemes are reducible to low-level auditory properties, such as pitch or timbre.

It may be helpful to round off this discussion by bringing its findings to bear on what is arguably the standard, programmatic way of dealing with claims about high-level perception: why think that these effects should be explained at the perceptual level at all, rather than at a postperceptual, cognitive level or at a preperceptual level of sensory registration?[13]

To meet this strategy, we must assess the balance of available empirical evidence and the trajectory of current theorizing on the matter. As such, we cannot expect knock-down arguments. However, I do not think there can be any doubt that the balance of evidence strongly favors the conclusion that these effects are indeed perceptual. Even though the capacities arise only as a result of the involvement of cognitive faculties, they are not themselves a matter of cognition. For instance, the desensitization to case described in section 2.2 occurs too early in visual processing for this strategy to be plausible; moreover, it occurs in a brain area (the fusiform gyrus), which, indisputably, is involved in distinctive forms of perceptual processing, even according to Burge and Block (i.e., face recognition). Finally, it bears remarking that the relevant experiments produce the same patterns of repetition suppression even under masking conditions, that is, conditions where subjects are unable to report the occurrence of the stimulus at all (c.f., Dehaene et al., 2001). These factors strongly suggest that we should not seek cognitive-level explanations for the phenomena in question.

On the other hand, the claim that the relevant effects are not preperceptual can be supported by insights into the systematic nature of reading deficits. Ironically, Fodor is hinting at the right question when he asks, rhetorically, in what sense we should "[suppose] illiterates to be visually incapacitated?" If we focus on developmental dyslexia, the answer is presumably that, in an important sense, they are not visually incapacitated at all. But that insight lends no support to the austerist view. In standard cases,[14] dyslexia has not been shown to robustly correlate with significant deficits in low-level visual processing (c.f., Castles, McLean, & McArthur, 2010, pp. 428–429). In these cases, subjects can *see* the lines, curves, and dots that make up letters as clearly as anyone else. To explain the nature of the deficit, we must postulate a more abstract level of perceptual processing: what they lack is the ability to perceptually group these lines, curves, and dots into graphemes with the speed and spontaneity that is required for fluent reading. While these graphemes are the product of a process that crucially involves cognitive capacities, as I have argued, the exercise of the resultant capacity itself is firmly located at the perceptual level.

Interestingly, Fodor's inferential model may indeed provide a plausible description of early stages of reading and perhaps even of certain kinds of developmental dyslexia. But as we can now

[13] Many thanks to an anonymous referee for pressing me on this point. Burge himself appears to invoke a strategy of this form in 2014, pp. 582–583.

[14] That is, cases that are explained by preexisting deficits in phonological processing. As Castles and Friedman (2014) point out, it is plausible that the "phonological deficit hypothesis" is not the only game in town and that there may be several distinct causal pathways that can manifest in reading deficits of this sort. This is a matter of bleeding-edge empirical inquiry (c.f., Ozernov-Palchik & Gaab, 2016). But it is sufficient for my purposes here that *some*—plausibly most—cases of developmental dyslexia are explained without invoking deficits either in preperceptual visual processing or in more general cognitive factors, such as attention or memory.

see, that description only serves to underscore how different we are as fluent readers. Becoming a fluent reader involves the acquisition of a new range of perceptual attributives. This process involves, at an early stage, the acquisition of new concepts—concepts of graphemes—which are enmeshed in inferential routines supported by high-level cognitive capacities, such as attention and memory. This, in itself, should not be a cause of concern or even surprise, insofar as it is precisely part of what we would mean by saying that the capacities in question are acquired by learning. But at some point, the inferential routine is dropped, and the relevant stimuli directly trigger the appropriate perceptual response. We have now acquired a new range of perceptual attributives, namely graphemes.

4 | PERCEPTUAL EXPANSION UNDER COGNITIVE GUIDANCE: A MATTER OF "COGNITIVE PENETRATION"?

These reflections can provide a new inroad into current philosophical debates about the reach of perception. True, they provide no direct support for claims that we can perceive everyday objects such as teacups and toasters.[15] In its own right, this is of no great concern to me—I would be happy to occupy a middle ground in a debate that has all too often been polarized between two extremes. But they do, I think, conclusively demonstrate, on empirical rather than phenomenological grounds, that austerists are wrong to suppose that we are not in possession of any high-level, culture-specific perceptual attributives. Moreover, paying close attention to the way that we come into possession of these perceptual attributives can help block argumentative strategies frequently deployed by austerists: in particular, the involvement of cognitive capacities at the acquisition stage does not settle the question of whether the resulting capacities are perceptual.

However, many contemporary philosophical discussions of the reach of perception are cast in terms of the concept of *cognitive penetration.* This way of framing the matter has arguably contributed to raising the evidential bar for claims about high-level perception. In this final section, I aim to show that there is reason to think that cognitive penetration may not be the most fruitful heading under which to discuss the issues that arise from this paper and that, in general, we would do well to separate these strands of the discussion.

Here, for instance, is how Fiona Macpherson approaches the question of cognitive penetration. Assume that we hold fixed the object of perception, the viewing conditions, the state of the sensory organs, as well as the location of attentional focus; now, "if it is possible for two subjects in these conditions to have different perceptual experiences (different in respect of phenomenal character and content) on account of the differing states of their cognitive systems, […] then cognitive penetration is possible" (Macpherson, 2012 p. 28; see also Stokes, 2013, pp. 647–648). On such a minimal definition, however, speech perception would quite clearly show that cognitive penetration is not only possible but, in fact, perfectly commonplace. To return to an example given above, even if we hold fixed these factors, monolingual speakers of English have different perceptual experiences of aspirated /t/s than do fluent speakers of Spanish. This difference is explained at the cognitive level, in terms of which languages they know.

Presumably, then, this is not what philosophers who take an interest in cognitive penetration usually have in mind. What is supposed to be controversial about the cognitive penetration thesis is

[15] Though note that the results also may not be specifically restricted to language processing: see Hannagan *et al.*, 2015, for discussion of the possibility that similar learning processes may be capable of tuning the ventral occipitotemporal cortex into a wider range of visual "symbols," including numerals.

that it seems to jeopardize the neutrality of experience (for an overview, see Silins, 2016). If our faculties or belief or desire are constantly busying themselves with interfering in occurrent perceptual processing, then it is not clear that simple observations can provide an epistemological foundation for justifying empirical beliefs in the way that we commonly take for granted. Instead, what we see (hear, etc.) is deeply influenced by what we believe (feel, etc.). This is a real concern.[16] But however surprising and even startling these findings from recent psychological and neuroscientific inquiries into language processing may be, I see no grounds for thinking that they raise this sort of specter. The outputs of language processing in the auditory and visual modalities are not idiosyncratic, but meet every sensible and relevant test of objective ascertainability. Meanwhile, to imagine a situation in which, say, the perceptual output of speech processing were "more true" to the actual acoustic signal than it is, is essentially just to imagine the devastating (but very real) condition known as pure word deafness (c.f., Krival, 2011), where subjects hear the speech signal as pure sound with no discernible phonemic structure.

These reflections suggest that cognitive penetration is not always the most fruitful heading under which to discuss the question of high-level perception.[17] "Perceptual learning" would be a handy term, were it not for the fact that some of the psychologists who deploy it do not carefully distinguish these results from claims about cognitive penetration, according to which our highly idiosyncratic doxastic or affective states are continuously imposing themselves on real-time perceptual processing (c.f., Goldstone, de Leeuw, & Landy, 2015). In contrast to these quite radical claims, I have argued for the more moderate thesis that there exist processes straddling the perception–cognition interface that can yield products that add to the stock of perceptual attributives *over time.*[18] But even such a moderate argument stands in opposition to more austere views of the reach of perception, such as those defended by Burge and Block, according to which our perceptual repertoires must remain tied to a largely fixed stock of low-level attributives, only augmented by a small range of somewhat higher-level attributives (e.g., integrated body, food, shelter, aspects of face perception, and so on). Just how far this template can stretch—and whether it could begin to impinge on the territory more regularly asserted by perceptual liberals—is a strictly empirical matter, on which we may patiently have to await further evidence. Meanwhile, it is important to note that there is an important sense in which my conclusions leave the operations of our perceptual systems autonomous, in precisely the way that should matter to austerists:[19] so far as my arguments give us any reason to believe, it is still true that, *at any given time*, our perceptual systems operate in relative independence of our cognitive faculties and perform computations only with such attributives as it has in stock.

ACKNOWLEDGEMENTS

Many thanks to Kathleen Akins, Holly K. Andersen, Josh Armstrong, Ned Block, Christopher Gauker, Steven Gross, Martin Hahn, Tristram Oliver-Skuse, Jessica Pepp, David Pereplyotchik, and

[16] For a systematic critical review of psychological studies claiming to show that higher mental states and processes, such as moods and assessments of task difficulty, might influence the contents of perception, see Firestone & Scholl, 2016.

[17] Brogaard & Chomanski, 2015, voice a similar concern, though drawing on very different kinds of arguments.

[18] Churchland's "diachronic cognitive penetration" (Churchland, 1988) could also have been a useful coinage but for the fact that he persists in framing the issue—unhelpfully in my view—in terms of the theory-ladenness of observation.

[19] It is worth noting that cognitive penetration is not discussed as such in Burge's book but is quite central to the way he frames the issue in Burge (2014, pp. 582–583). I am hopeful that this shift of target is indicative of a growing awareness that the real threat to his picture of the mind is not, *per se*, the acquisition of new high-level perceptual attributives (where these result from empirically explicable learning processes crossing the perception–cognition divide) but rather claims about the real-time imposition of cognitive states on perceptual processing.

Markos Valaris, as well as audiences at the Society for Philosophy and Psychology (June 2014) and the University of Rochester (April 2015) for discussions that led to significant improvements of this paper. A special note of gratitude is due to Rob Stainton, who first encouraged me to think more closely about the philosophical challenges arising from language perception.

REFERENCES

Baker, S. A., Idsardi, W. J., Golinkoff, R. M., & Petitto, L.-A. (2005). The perception of handshapes in American Sign Language. *Memory & Cognition*, *33*, 887–904.

Bayne, T. (2009). Perception and the reach of phenomenal content. *The Philosophical Quarterly*, *59*, 385–404.

Begby, E. (2011). Review of Tyler Burge, Origins of objectivity. *Notre Dame Philosophical Reviews*, *2011*(2).

Best, C. T., Mathur, G., Miranda, K. A., & Lillo-Martin, D. (2010). Effects of sign language experience on categorical perception of dynamic ASL pseudosigns. *Attention, Perception, & Psychophysics*, *72*, 747–762.

Block, N. (2014). Seeing-as in the light of vision science. *Philosophy and Phenomenological Research*, *89*, 560–572.

Bolger, P., Borgwaldt, S. R., & Jakab, E. (2009). Letter and grapheme perception in English and Dutch. *Written Language & Literacy*, *12*, 116–139.

Brogaard, B. (2013). Do we perceive natural kind properties? *Philosophical Studies*, *162*, 35–42.

Brogaard, B., & Chomanski, B. (2015). Cognitive penetrability and high-level properties in perception: Unrelated phenomena? *Pacific Philosophical Quarterly*, *94*, 469–486.

Burge, T. (2010). *Origins of objectivity*. Oxford: Oxford University Press.

Burge, T. (2014). Reply to Block: Adaptation and the upper border of perception. *Philosophy and Phenomenological Research*, *89*, 573–583.

Casserly, E. D., & Pisoni, D. B. (2010). Speech perception and production. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*, 629–647.

Castles, A., & Friedman, N. (2014). Developmental dyslexia and the phonological deficit hypothesis. *Mind & Language*, *29*, 270–285.

Castles, A., McLean, G. M. T., & McArthur, G. (2010). Dyslexia (neuropsychological). *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*, 426–432.

Churchland, P. M. (1988). Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science*, *55*, 167–187.

Dehaene, S. (2009). *Reading in the brain: The science and evolution of a human invention*. New York, NY: Viking Press.

Dehaene, S., Naccache, L., Cohen, L., Le Bihan, D., Mangin, J.-F., Poline, J.-B., & Rivière, D. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, *4*, 752–758.

Dretske, F. (1995). *Naturalizing the mind*. Cambridge, MA: The MIT Press.

Emmorey, K., McCullough, S., & Brentari, D. (2003). Categorical perception in American Sign Language. *Language and Cognitive Processes*, *18*, 21–45.

Fernandez, E. M., & Smith Cairns, H. (2010). *Fundamentals of psycholinguistics*. Oxford: Wiley-Blackwell.

Firestone, C., & Scholl, B. (2016). Cognition does not affect perception: Evaluating the evidence for 'top-down' effects. *Behavioral and Brain Sciences*, *39*, 1–77.

Fish, W. (2013). High-level properties and visual experience. *Philosophical Studies*, *162*, 43–55.

Fodor, J. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: The MIT Press.

Fodor, J. (1988). A reply to Churchland's 'Perceptual plasticity and theoretical neutrality'. *Philosophy of Science*, *55*, 188–198.

Fodor, J., Bever, T. G., & Garrett, M. F. (1974). *The psychology of language: An introduction to psycholinguistics and generative grammar*. New York, NY: McGraw Hill.

Frith, U. (1985). Beneath the surface of developmental dyslexia. In K. Patterson, J. Marshall, & M. Coltheart (Eds.), *Surface dyslexia, neuropsychological and cognitive studies of phonological reading*. London: Erlbaum.

Gibson, E. J., Pick, A., Osser, V., & Hammond, V. (1962). The role of grapheme-phoneme correspondence in the perception of words. *The American Journal of Psychology*, *75*, 554–570.

Goldstone, R. L., de Leeuw, J. R., & Landy, D. H. (2015). Fitting perception in and to cognition. *Cognition*, *135*, 24–29.

Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*, 69–78.

Goto, H. (1971). Auditory perception by normal Japanese adults of the sounds 'l' and 'r'. *Neuropsychologia*, *9*, 317–323.

Hannagan, T., Amedi, A., Cohen, L., Dehaene-Lambertz, G., & Dehaene, S. (2015). Origins of the specialization for letters and numbers in ventral occipitotemporal cortex. *Trends in Cognitive Sciences*, *19*, 374–382.

Hawley, K., & Macpherson, F. (Eds.) (2011). *The admissible contents of experience*. Oxford: Wiley-Blackwell.

Jusczyk, P. W. (2003). The role of speech perception capacities in early language development. In M. T. Banich & M. Mack (Eds.), *Mind, brain, and language: Multidisciplinary perspectives*. Mahwah, NJ: Lawrence Erlbaum Associates, Publishers.

Kluender, K. R., & Kiefte, M. (2006). Speech perception within a biologically realistic information-theoretic framework. In M. Traxler & M. Gernsbacher (Eds.), *Handbook of psycholinguistics* (2nd ed.). Amsterdam: Elsevier.

Krival, K. (2011). Pure word deafness. In J. S. Kreuzer, J. DeLuca, & B. Caplan (Eds.), *Encyclopedia of clinical neuropsychology*. New York, NY: Springer.

Liberman, A. M., Cooper, F. S., Shankweiler, D. P., & Studdert-Kennedy, M. (1967). Perception of the speech code. *Psychological Review*, *74*, 431–461.

Lupker, S. J., Acha, J., Davis, C. J., & Perea, V. (2012). An investigation of the role of grapheme units in word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *38*, 1491–1516.

Macpherson, F. (2012). Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research*, *84*, 24–62.

Marinus, E., & de Jong, V. (2011). Dyslexic and typical-reading children use vowel digraphs as perceptual units in reading. *The Quarterly Journal of Experimental Psychology*, *64*, 504–516.

Masrour, F. (2011). Is perceptual phenomenology thin? *Philosophy and Phenomenological Research*, *83*, 366–397.

McCandliss, B. D., Cohen, L., & Dehaene, S. (2003). The visual word form area: Expertise for reading in the fusiform gyrus. *Trends in Cognitive Science*, *7*, 293–299.

O'Callaghan, C. (2011). Against hearing meanings. *The Philosophical Quarterly*, *61*, 783–807.

O'Callaghan, C. (2014). Speech perception: Empirical and theoretical considerations. In *Stanford Encyclopedia of Philosophy*. Stanford, CA: Center for the Study of Language and Information.

Ozernov-Palchik, O., & Gaab, N. (2016). Tackling the 'dyslexia paradox': reading brain and behavior for early markers of developmental dyslexia. *Wiley Interdisciplinary Reviews: Cognitive Science*, *7*, 156–176.

Pegado, F., Comerlato, E., Ventura, F., Jobert, A., Nakamura, K., Buiatti, M., … Dehaene, S. (2014). Timing the impact of literacy on visual processing. *Proceedings of the National Academy of Sciences*, *111*, E5233–E5242.

Pinker, S., & Jackendoff, V. (2005). The faculty of language: What's special about it? *Cognition*, *95*, 201–236.

Polk, T. A., & Farah, M. J. (2002). Functional MRI evidence for an abstract, not perceptual, word-form area. *Journal of Experimental Psychology: General*, *131*, 65–72.

Prinz, J. (2005). A neurofunctional theory of consciousness. In A. Brook & K. Akins (Eds.), *Cognition and the brain: The philosophy and neuroscience movement*. Cambridge: Cambridge University Press.

Pylyshyn, Z. (1984). *Computation and cognition*. Cambridge, MA: The MIT Press.

Pylyshyn, Z. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, *22*, 341–365.

Reid, T. (1764). *An inquiry into the human mind on the principles of common sense*. In W. Hamilton (Ed.), *The works of Thomas Reid* (Vol. *1*). Edinburgh: MacLachlan and Stewart, 1872.

Reid, T. (1785). *Essays on the intellectual powers of man*. In W. Hamilton (Ed.), *The works of Thomas Reid* (Vol. 1). Edinburgh: MacLachlan and Stewart, 1872.

Reiland, I. (2014). On experiencing high-level properties. *American Philosophical Quarterly*, *51*, 177–187.

Rey, A., Ziegler, J. C., & Jacobs, A. M. (2000). Graphemes are perceptual reading units. *Cognition*, *75*, B1–12.

Sheldon, A., & Strange, W. (1982). The acquisition of /r/and/l/ by Japanese learners of English: Evidence that speech production can precede speech perception. *Applied Psycholinguistics*, *3*, 243–261.

Siegel, S. (2006). Which properties are represented in perception? In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual experience*. Oxford: Oxford University Press.

Siegel, S. (2010). *The contents of visual experience*. Oxford: Oxford University Press.

Silins, N. (2016). Cognitive penetration and the epistemology of perception. *Philosophy Compass*, *11*, 24–42.

Stokes, D. (2013). Cognitive penetrability of perception. *Philosophy Compass*, *8*, 646–663.

Trout, J. D. (2001). The biological basis of speech: What to infer from talking to the animals. *Psychological Review*, *108*, 523–549.

Tye, M. (1995). *Ten problems of consciousness*. Cambridge, MA: MIT Press.

Watson, M. R., Blair, M. R., Kozik, P., Akins, K. A., & Enns, J. T. (2012). Grapheme-color synesthesia benefits rule-based category learning. *Consciousness and Cognition*, *21*, 1533–1540.

Werker, J. F., & Tees, R. C. (2002). Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development*, *25*, 121–135.

Zeimbekis, J., & Raftopoulos, A. (Eds.) (2015). *The cognitive penetrability of perception: New philosophical perspectives*. Oxford: Oxford University Press.

How to cite this article: Begby E. Perceptual expansion under cognitive guidance: Lessons from language processing. *Mind Lang*. 2017;32:564–578. https://doi.org/10.1111/mila.12159