Language and Linguistics Compass 5/2 (2011): 78–94, 10.1111/j.1749-818x.2010.00264.x

Functionally Localizing Language-Sensitive Regions in Individual Subjects With fMRI: A Reply to Grodzinsky's Critique of Fedorenko and Kanwisher (2009)

Evelina Fedorenko* and Nancy Kanwisher
McGovern Institute for Brain Research, MIT

Abstract

In Fedorenko and Kanwisher (F&K 2009), we argued that defining regions of interest functionally in individual subjects may lead to a clearer picture of the functional architecture of the language system because it affords higher sensitivity and selectivity. Grodzinsky (2010) takes issue with two aspects of the F&K paper. First, he argues that the picture of the neural basis of language that has emerged from previous work is not as murky as F&K argue, with the implication that perhaps a new method is not needed. And second, he raises some concerns with the individual-subjects functional localization approach and argues instead for the use of probabilistic cytoarchitectonic maps (e.g., Amunts et al. 1999). In the current manuscript, we respond to both of Grodzinsky's points. Regarding the first point, we stand by our assessment of the brain basis of language literature, and argue that many core questions related to functional specialization remain unanswered. We explain why we remain unconvinced by Grodzinsky's example of Broca's area (or a portion of it) as a region with a well-understood functional profile. Regarding the concerns that Grodzinsky raises with respect to the use of functional localizers in investigating the language system, we respond to the main comments and refer the reader to Fedorenko et al. (2010) where the remaining concerns are addressed (with both data and argumentation). Finally, we point out that the probabilistic cytoarchitectonic maps that Grodzinsky advocates have some of the same limitations as other group-based methods. We conclude that the individual-subjects functional localization approach, advocated in F&K, holds promise for better understanding the brain basis of language by enabling a detailed characterization of language-sensitive regions with respect to their role in both linguistic and non-linguistic processes.

1. Introduction

In Fedorenko and Kanwisher (2009; henceforth F&K 2009), we made an argument for the need for a new approach to the study of the neural basis of language using fMRI. In particular, we argued that many fundamental questions about the brain regions supporting language production/comprehension processes remain unanswered, and we advocated a method where regions of interest are defined functionally in individual subjects. This method circumvents the problem of anatomical variability (which leads to imperfect alignment of functional activations), and has been highly effective in characterizing brain regions supporting high-level visual processes (see e.g., Kanwisher 2010, for a recent review) and those underlying social cognition (e.g., Saxe and Powell 2006). We therefore hypothesized that this method may also help reveal a clearer picture of the functional organization of the language system. Grodzinsky (2010) questions F&K's assessment of the literature, arguing that the emerging picture of the brain basis of language is clearer than suggested by F&K. As a result, in his opinion, a new method is not needed.

Furthermore, Grodzinsky raises a number of concerns with the individual-subjects functional localization approach. He instead advocates the use of anatomical regions of interest based on probabilistic cytoarchitectonic maps (e.g., Amunts et al. 1999). In the current manuscript, we respond to both of these points, with a somewhat stronger focus on the second, methodological, point because it is the core of the F&K's (2009) paper. Moreover, the first argument has already received a lot of attention in the literature, which we allude to below.

Although the two arguments (about the current state of affairs in neurolinguistics, and about the appropriate methods to use in evaluating hypotheses about the brain basis of language) are intertwined throughout Grodzinsky's response, these issues are, to a large extent, orthogonal, and so we will treat them as such in our response.[1]

2. The Current State of Affairs With Respect to the Neural Basis of Language

Because (a) the original goals of the F&K paper did not include an extensive review of the brain basis of language literature, and (b) many of these issues have already received sufficient attention in the literature, we will only respond to a few key points.

2.1. CLARIFYING OUR POSITION

2.1.1. Relying on Meta-analyses is Not Critical for Arguing for a Murky Picture of the Brain Basis of Language

Grodzinsky points out a number of problems with the meta-analyses that we referred to in F&K (i.e., Lindenberg et al. 2007; Ojemann 1991; Vigneau et al. 2006). In particular, he argues that these (and other) meta-analyses may not accurately reflect the state of affairs because they often combine data from studies that vary along multiple dimensions. Consequently, some patterns of results that may be clear and consistent across similar individual studies will get lost when these studies are lumped together with less carefully controlled studies or studies examining superficially similar, but actually underlyingly different functional contrasts. We whole-heartedly agree with these points (discussed in sections 2 and 3 in Grodzinsky's response). However, relying on meta-analyses is not critical for the claim that we made, and instead served as a *shortcut*. We felt justified in using this shortcut because we did not think that many clear patterns of functional dissociations were hidden in the summary of the literature that we provided (see footnote 2 in F&K where we list a few of the most consistent findings with respect to the relationship between different brain regions and linguistic functions).

2.1.2. We are Not Making Any Arguments (Yet) About the Brain Basis of Language

In several instances throughout the paper, Grodzinsky incorrectly represents our position with respect to the neural basis of language. For example, on p. 616, he says '… has led FK to argue that no brain region is either necessary or sufficient for language processing, let alone syntax'. However, we have not actually made any claims about the necessity, sufficiency or functional specificity of any language-sensitive brain regions. What we have argued is the following (see Fedorenko et al. 2010, for additional discussion): although some linguistic tasks quite consistently activate particular brain regions, traditional group-based methods have so far largely failed to detect functional specificity that is replicable across studies and labs (see Section 3, for additional discussion). As a result, for any language-sensitive brain region, it is not yet known whether this region (a) selectively supports a particular linguistic process (e.g., processing speech sounds, or forming

dependencies between syntactic elements), or (b) selectively supports language (i.e., does not support any non-linguistic processes, such as, for example, general working memory, arithmetic processing, or musical processing). We believe that part of the difficulty in uncovering functional dissociations – among different aspects of language, or between language and non-linguistic processes – results from the tendency of group-based methods to underestimate selectivity.

2.2. BROCA'S AREA IS NOT A GOOD EXAMPLE OF CLARITY IN THE NEURAL BASIS OF LANGUAGE LITERATURE

In discussing the degree of murkiness/clarity in the picture of the neural basis of language that emerged from previous work, Grodzinsky focuses exclusively on Broca's area. Setting aside the fact that language is supported by a host of brain regions and, as stated above, in our opinion, no consensus exists on the role of *any* of these language-sensitive regions in language processing or on the extent to which they are specialized for language, even with respect to Broca's area, Grodzinsky's representation of the available evidence is incomplete.

At the core of Grodzinsky's argument for a 'stable and clear' (p. 605), picture of the neural basis of language is the claim that Broca's area [at least the Brodmann Area (BA) 45 portion of it] selectively supports syntactic processing. In talking about the function of Broca's area, Grodzinsky invokes the notion of 'syntactic movement'. This idea comes from Chomsky's (1965) proposal whereby linguistic structures in the form that they appear in the language (surface representations) are 'derived' from underlying core forms (deep representations) via a series of 'transformations' (that are also referred to as 'syntactic movement', e.g., Chomsky 1973; and many more recent proposals, such as Chomsky 1981, 1993; Pollard and Sag 1994). For example, in a passive structure, as in (1), 'Mary' is argued to 'move' from its base position as the object of the verb 'kissed' to the subject position, leaving a 'trace' behind.

(1) Mary_1 was kissed t_1 by John.

It is the process of linking the displaced element ['Mary' in (1)] and its trace [in the object position of 'kissed' in (1)] that Grodzinsky argues (a) leads to activity in Broca's area, and (b) is impaired in patients with damage to Broca's area. (It is worth noting that this is only one way to analyze these and other structures. For example, alternative linguistic accounts of the passive construction in (1) (e.g., Bresnan, 1982) proposed direct associations between verbs and their arguments.)

We here briefly outline three key problems with the proposal that Broca's area is selectively engaged in syntactic movement, and argue that because of the controversies that still surround left inferior frontal regions, Broca's area and its role in syntactic processing is not a good example of clarity in our understanding of the brain basis of language (see also a recent debate in TiCS between Grodzinsky and Santi 2008 and Willems and Hagoort 2009).

2.2.1. Problems With Grodzinsky's Proposal for the Selective Role of Broca's Area in Syntactic Movement

2.2.1.1. 'Consistency'/'regularity' does not equal 'specificity'. In order to argue that some brain region R is selectively engaged in some cognitive process X,[2] two conditions must be met: (1) tasks relying on cognitive process X must *consistently activate* region R, and (2)

tasks relying on other cognitive processes (e.g., Y or Z) must *not* activate region R. Most of the evidence that Grodzinsky discusses in his response to F&K and elsewhere (e.g., Grodzinsky and Santi 2008) are relevant to the first condition, i.e., establishing a consistent relationship between Broca's area and a particular kind of a linguistic contrast. For example, on p. 612, Grodzinsky says, '… the striking regularity above is no small matter, strongly suggesting that *even relatively gross stroke-induced lesions display functional specificity of a kind that the meta-analyses which FK reviewed completely missed*'; and then again, on the same page, '… clear effects are obtained, seemingly holding across modalities, tasks, materials and analytic methods and are strong enough to create a reasonably clear picture'. Grodzinsky's reasoning seems to be as follows: because a relatively narrow/subtle linguistic contrast (i.e., a contrast between structures that do vs. do not involve syntactic movement, or that differ in the distance between the linguistic elements which are dependent on one another) consistently activates a particular brain region (i.e., BA 45), this region is selectively engaged in the linguistic operation in question (i.e., syntactic movement). This reasoning is flawed, however. All that can be concluded based on this evidence is that BA45 is *consistently activated* by (and perhaps, is *critical* for, based on the evidence from aphasia) syntactic movement.[3]

2.2.1.2. Arguing for functional selectivity while ignoring evidence inconsistent with the selective role of Broca's area in syntactic movement is problematic. Although Grodzinsky's group does contrast syntactic movement manipulations with other linguistic manipulations, these contrasts are far from sufficient for concluding that Broca's area (or BA45) is selectively engaged in syntactic movement.[4] In particular, much evidence exists suggesting that left inferior frontal regions support non-syntactic aspects of language, such as phonological processing (e.g., Blumstein et al. 2005; Myers et al. 2009) or lexico-semantic processing (Hagoort et al. 2004, 2009; Rodd et al. 2005; Schnur et al. 2009), as well as non-linguistic cognitive processes, such as arithmetic processing, working memory/cognitive control processes, or musical processing (e.g., Dehaene et al. 1999; Duncan 2001; Levitin and Menon 2003; Maess et al. 2001; Owen et al. 2005). As discussed above, given the tendency of group-based methods to underestimate functional specificity, it remains an open question whether some portions of the left inferior frontal cortex are truly multifunctional, or whether subregions exist within those regions that support distinct cognitive processes. Critically, however, this kind of evidence is of utmost importance for claims about functional specificity.

Grodzinsky acknowledges some of this evidence at the beginning of section 7. He raises a possibility that some of these results would not 'withstand careful scrutiny' (p. 616). We are very sympathetic to this objection: there is no shortage of neuroimaging studies of language that are not carefully controlled, not properly analyzed, and/or would not replicate on a new group of subjects. However, we strongly disagree with the statement that Grodzinsky makes next, 'But even if some [studies] do [withstand careful scrutiny], it is important to realize that the current debate remains unaffected'. If we are correctly interpreting what debate Grodzinsky is referring to here (i.e., the debate about the functional profile of Broca's area), then the debate is *hugely* affected by findings of this kind.[5] For example, if one conclusively shows that BA45 is activated by both (a) syntactic movement, but also (b) general working memory, then the hypothesis about selective engagement of this region in syntactic movement can be rejected, and new hypotheses that can account for all the relevant findings will need to be generated (see Willems and Hagoort 2009, for additional discussion).

Of course, there are different research strategies: some may choose to study all the cognitive processes that have been argued/shown to activate a particular region and make inferences based on the patterns that emerge with respect to a host of different tasks/stimuli, while others may choose to focus on one particular process that consistently activates a particular region and try to carefully characterize the region's engagement in this process. Both of these strategies are important for ultimately understanding the computations that the region performs and the representations that it stores and/or manipulates. However, if one chooses the latter strategy (as Grodzinsky is doing by characterizing the engagement of BA45 in syntactic processing), then one is not justified in making claims about functional specificity of the region, especially in the presence of substantial evidence inconsistent with such claims. In the abstract of the paper, Grodzinsky says, 'when the neurolinguistic landscape is examined with the right linguistic spectacles, the emerging picture – while intriguingly complex – is not murky, but rather, stable and clear' (p. 605). Surely, however, these 'linguistic spectacles' should not make one selectively focus on just the evidence consistent with a particular hypothesis and ignore other, highly relevant, evidence.

In summary, in order to argue for a *selective* role of Broca's area in syntactic processing, one needs to systematically compare syntactic tasks to the multitude of other linguistic and non-linguistic tasks that have been shown to activate cortex in/around Broca's area (see e.g., Willems and Hagoort 2009 or Fedorenko et al. forthcoming, for additional discussion). Failure to do so amounts to succumbing to confirmation bias (e.g., Nickerson 1998; Wason 1960). Furthermore, as we have argued in F&K (2009) and in Fedorenko et al. (2010), it is critical to perform these analyses in individual subjects because group analyses may obscure specificity, especially in cases of adjacent functionally distinct regions. (Indeed, it is possible that the individual subject method we advocate might produce just the pattern of results Grodzinsky's hypothesis predicts: selective activation of a portion of Broca's area only by syntactic movement manipulations and not by any other linguistic or non-linguistic manipulations.)

2.2.1.3. Grodzinsky's characterization of behavioral psycholinguistic evidence is flawed. With respect to behavioral psycholinguistic work, Grodzinsky selectively cites studies that he interprets as being consistent with his theoretical position and fails to cite other, highly relevant, studies that either undermine the findings he refers to or provide alternative explanations for the empirical phenomena in question.

For example, in his response to F&K, Grodzinsky cites a study by Nicol and Swinney (1989) as evidence for the psychological reality of syntactic movement (interpreted by him therefore as convergent with his neuroimaging findings). These researchers used a cross-modal lexical priming paradigm where participants are presented with auditory sentence stimuli and visual probes – presented at different times relative to the auditory sentence stimulus – on which a lexical decision has to be made. The idea was to test whether facilitation (priming) effects could be observed when a word semantically related to the displaced element (or *filler*) is presented at the 'trace' (or *gap*) position. Nicol and Swinney indeed observed such an effect, as Grodzinsky reports. However, McKoon and colleagues (McKoon and Ratcliff 1994; McKoon et al. 1996[6]) later demonstrated that the effects reported by Nicol and Swinney were due to a confound in the experimental design. In particular, the words for which priming was observed were also possible objects/patients of the verb after which they were presented. Controlling for this confound eliminated the priming effects.

Furthermore, although Grodzinsky does acknowledge possible alternative interpretations of the syntactic contrasts he investigates (e.g., on p. 611 he says, 'You may pick your favorite explanation for these results'), he then proceeds to talk about the relevant studies as evidence for a *syntactic* account of the role of Broca's area, even though many explanations of the contrasts in question rely on not necessarily syntax- or even language-specific factors (e.g., Fedorenko et al. 2006, 2007; Gennari and MacDonald 2009; Gibson 1998; Gordon et al. 2002; Lewis et al. 2006; McElree et al. 2003). Such persistent labeling of empirical effects that are consistent with many alternative hypotheses in terms of a specific theoretical framework is misleading and may obscure generalizations that could become more apparent when more theory-neutral terms are used to describe the findings.

3. fROIs Defined in Individual Subjects vs. Probabilistic Cytoarchitectonic Maps

We completely agree with Grodzinsky that because the individual-subjects functional localization approach 'marks a departure from current practice' (p. 606), it needs to be carefully evaluated and shown to be at least as good as, or better than, currently used methods. We do precisely that in our recent work. In Fedorenko et al. (2010), we develop and validate one possible localizer for defining language-sensitive brain regions. In particular, using a contrast between sentences and pronounceable non-words, we identify a set of brain regions that are (i) present in the vast majority of individual subjects, (ii) show replicable patterns of activity within subjects and across subject groups, (iii) respond similarly to linguistic stimuli presented visually vs. auditorily, and (iv) respond similarly across different tasks (passive reading vs. reading with a memory probe). These regions include the classically implicated regions on the lateral surfaces of the left frontal lobe and left temporal/parietal lobes, as well as a few additional regions (see http://web.mit.edu/evelina9/www/funcloc.html for all the tools that we developed for performing subject-specific fROI analyses). Furthermore, in Fedorenko et al. (forthcoming), we demonstrate that this method yields higher functional selectivity compared to traditional, group-based, methods. Finally, in Nieto-Castañon et al. (forthcoming), we provide a formal discussion – supported by a series of simulations – of how the individual-subjects fROI approach compares to traditional (group-level voxel- and ROI-based) approaches in terms of sensitivity and selectivity.

In the remainder of this section, we (1) clarify our position on several issues that were not accurately characterized in Grodzinsky's response, (2) respond to the concerns raised by Grodzinsky with respect to the individual-subjects fROI method, and (3) outline some concerns with the probabilistic cytoarchitectonic maps that Grodzinsky advocates. In order to better situate the current debate, we start by laying out the space of possible fMRI analysis strategies.[7] We summarize these possibilities in Figure 1.

First, brain activity can be examined via statistical tests computed on each voxel (a three-dimensional pixel) in the brain (*whole-brain analyses*) either in an individual subject or across subjects aligned in a common space. Examining individual subject activation maps can be highly informative, but in order to claim that a certain pattern of activity is characteristic of the population more generally, activation patterns need to be compared across individuals so that inferential statistical tests can be performed. In order to compare activations across individual brains, brains need to be aligned in a common space, so that correspondence can be established between any given point in one brain and a similar point in every other brain. (This process of lining up brains together in a common space is necessarily imperfect, as discussed in F&K (2009), due to inter-subject anatomical variability.) Then, in each unit of analysis – in this case, a voxel – a statistical test is

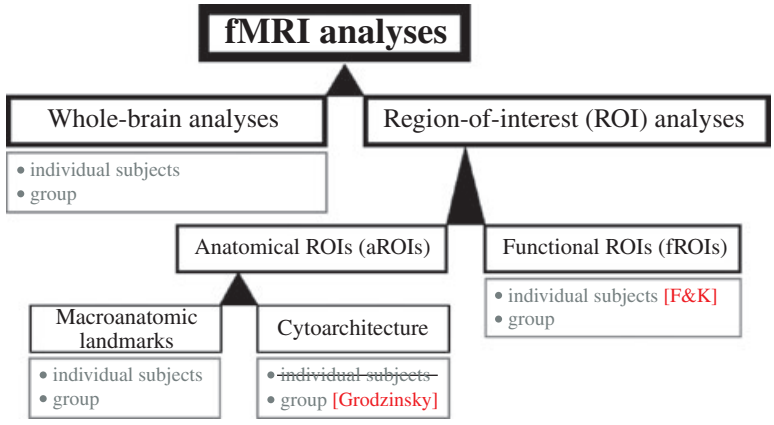

Fig 1. The space of possible analyses in functional MRI. F&K's and Grodzinsky's positions are indicated in red. (Defining cytoarchitecture-based ROIs in individual-subjects is under active exploration but not widely available; hence this option is crossed out.)

performed to determine whether the data patterns are similar across subjects (e.g., a *t*-test in each voxel can be used to determine whether one condition elicited a reliably stronger response than another condition).

In contrast to whole-brain analyses, *ROI analyses* focus on specific brain regions. Whereas whole-brain analyses are aimed at identifying a region/set of regions in the brain that is/are sensitive to a particular experimental manipulation, ROI-based analyses are aimed at testing hypotheses about particular brain regions, determined *a priori*. A region of interest can be defined based purely on the anatomy (anatomical ROIs, aROIs) or on some functional contrast (functional ROIs, fROIs). In defining anatomical ROIs, one can rely on macroanatomic landmarks (gyri and sulci; e.g., inferior frontal gyrus, or superior temporal sulcus; or subcortical structures, like the amygdala or the hippocampus) or on the microanatomy (cytoarchitectonic areas; e.g., BAs). Macroanatomically defined anatomical ROIs are commonly used in the literature (these are typically defined on a template brain; cf. Nieto-Castanon et al. 2003, for a demonstration of improved sensitivity/selectivity in individually defined aROIs). However, cytoarchitecture, not macroanatomy, has been shown to correspond to function (e.g., Iwamura et al. 1983; Matelli et al. 1991; Rozzi et al. 2008). Unfortunately, using even high-resolution fMRI, it is rarely possible to see below the level of macro-architecture on the cortical surface of any given brain (with a few exceptions; e.g., Clark et al. 1992). (Hence this possibility is crossed out in Figure 1 above.) Until a few years ago, only macroanatomically defined aROIs were possible (modulo the very crude attempts to estimate the locations of cytoarchitectonic zones based on macroanatomy). However, as Grodzinsky points out, heroic work by Zilles, Amunts and colleagues (e.g., Amunts et al. 1999, 2000, 2005; Caspers et al. 2006; Morosan et al. 2001) – known as the Julich Brain Mapping Project – has made it possible to estimate the locations of cytoarchitectonic areas at the group level by projecting the locations of cytoarchitectonic areas – defined in a set of 10 post-mortem brains – into a normalized space (used for aligning brains together to perform group analyses, as discussed above). These maps, Grodzinsky argues, present a perfect solution for characterizing different brain regions.

Finally, a region of interest can be defined functionally, by using a contrast aimed at the cognitive process of interest (for example, a contrast between faces and objects is used

to define face-selective regions). Similar to aROIs, functional ROIs can be defined using a group-level map (as done quite commonly in the literature; e.g., Kuperberg et al. 2003) or using individual activation maps (as done in a handful of papers[8] in the previous literature; Ben-Shachar et al. 2004; January et al. 2009; Hickok et al. 2009; see Fedorenko et al. 2010, for a brief discussion). In F&K (and also in Fedorenko et al. 2010, forthcoming), we argued for the use of standardized functional localizers to be run in individual subjects to identify candidate language regions. Developing and validating such localizers will enable more meaningful averaging in group-level analyses. In particular, functional localizers enable averaging data from corresponding functional regions across subjects, instead of corresponding locations in stereotaxic space, which may differ functionally due to inter-subject differences in the anatomy.

3.1. CLARIFYING OUR POSITION

Before discussing the arguments raised by Grodzinsky about fMRI analysis methods in his response, we want to clarify two issues with respect to our position on the methods that should be used in investigating the brain basis of language.

First, in contrast to Grodzinsky's claims (e.g., on p. 617, 'it is hard to say why it, and nothing else, is the right method to discover the language-relevant fROI'), we do not think (and have never claimed) that the particular localizer contrast that we chose to use in our initial explorations, or the individual-subjects functional localization approach more generally, is the *only* viable approach and/or that other analysis methods are not useful. What we have argued instead is that functionally defining ROIs in individual subjects will lead to higher sensitivity (an ability to detect an effect when it is present) and higher selectivity (an ability to distinguish among conditions if these conditions indeed elicit differential responses). As a result, we suggested that for studying questions of functional specificity, individual-subject fROI analysis methods may be better suited than group-based methods. As we discuss in Fedorenko et al. (2010), fROI analyses should always be supplemented by individual-subject whole-brain analyses, and, in some cases, by traditional whole-brain (random effects) group analyses.

And second, in contrast to Grodzinsky's implication (e.g., on p. 606, 'FK propose a fresh start'), we have never argued for ignoring the large body of evidence that is a result of several decades of neuroimaging and lesion-based work on language. Instead, we argue that some results – especially results that are interpreted as evidence of overlapping brain structures supporting, e.g., two aspects of language, or some aspect of language and some non-linguistic cognitive function – should be interpreted with caution because group analyses are bound to overestimate overlap and underestimate specificity.

3.2. POTENTIAL CONCERNS WITH USING INDIVIDUALLY DEFINED FROIS

Grodzinsky focuses on one main concern with respect to the use of individually defined fROIs: namely, the choice of the functional contrast. We respond to this concern next (see Fedorenko et al. 2010, forthcoming, for additional discussion of this and other concerns; see the debate between Saxe et al. 2006 and Friston et al. 2006, about the use of fROIs more generally).

Grodzinsky summarizes the crux of this issue on p. 617, 'unless we understand the character of the functional contrast used as a delimiter, and the rationale for its significance, it is hard to say why it, and nothing else, is the right method to discover the language-relevant fROI'.[9] We agree that the choice of the contrast for a functional localizer

is crucial, and we think Grodzinsky is correct to raise this concern. A research program based on a localizer that does not make sense would be of limited use. For that reason, we have thought hard about how to harness the power of the individual-subjects functional localization approach for language in a fashion that avoids this problem. We have two responses to Grodzinsky's concern about the 'grain' of functional localizers.[10]

First, in initial attempts to apply the individual-subjects functional localization approach to language (or any other new domain), there are different strategies one might adopt. One strategy is to focus on a particular aspect of language and try to develop a way to functionally define regions supporting this one narrow cognitive process. Another strategy is to examine a superset of language-sensitive regions. To do so, a functionally broader contrast is needed. We adopted the latter strategy in our initial efforts and developed our first language localizer based on the contrast between sentences and lists of pronounceable non-words (targeting regions supporting word- and sentence-level processes). Adopting such a strategy by no means implies that we do not acknowledge structure in the language system below the level of the localizer contrast (cf. Grodzinsky, p. 613). In our current work, we are investigating potential dissociations within the language system (e.g., between lexico-semantic and syntactic processing, or between different aspects of syntax). We may discover that different language-sensitive regions are more strongly, or perhaps even selectively, engaged in a particular aspect of language, as some previous studies have suggested. We may also discover that subregions *within* our fROIs show distinct functional signatures (underlying the importance of performing whole-brain analyses in individual subjects, in addition to examining the responses of fROIs to the experimental conditions). In that case, we will divide our fROIs into subregions and treat these separately in subsequent work. (Note that experiments aimed at testing functional dissociations within the language system can be used not only to examine the responses of the fROIs defined using our main localizer contrast, but also to develop new localizers targeting narrower aspects of language.[11]). If another contrast identifies additional regions in future experiments or subregions of those we advocate now, robustly enough to serve as individual-subject localizers, we can add these regions to our set of fROIs. It is only through such a multipronged effort that the field will be able to solve the chicken-and-egg problem of developing functional localizers that enable us to carve the language system at its joints without knowing the locations of those joints in advance.

Second, our initial attempts to develop ways to separately localize lexical processing (by using the WordList > NonwordList contrast) and syntactic processing (by using the Jabberwocky > NonwordList contrast) were not successful (see Appendix D in Fedorenko et al. 2010). However, as discussed above, we are currently evaluating other, more sophisticated, functional contrasts specifically targeting these and other aspects of linguistic processing. By making all the tools we developed – for performing individual subject analyses – publicly available from our website (http://web.mit.edu/evelina9/www/funcloc.html), we hope to encourage other researchers to work on developing functional localizers for particular aspects of language, given the advantages that the individual subjects functional localization approach affords. For example, if Grodzinsky develops a localizer for syntactic processing that works in individual subjects, we would be delighted to use it in our work. And as we develop additional localizers for various aspects of language, we will make them available from our website.

We explicitly say in Fedorenko et al. (2010) that the first localizer we developed is, by no means, the only possible localizer for language. However, the fact that the contrast we chose identifies all the key language-sensitive regions previously implicated in the

literature, and that the regions that appear systematically across subjects are functionally stable and exhibit properties of high-level language regions (e.g., a similar response to linguistic stimuli presented visually vs. auditorily) suggests that this localizer – targeting regions engaged in lexical- and sentence-level processing – is a reasonable place to start.

3.3. PROBABILISTIC CYTOARCHITECTONIC MAPS: ADVANTAGES AND POTENTIAL CONCERNS

From our perspective, the key advantage of the probabilistic cytoarchitectonic maps over whole-brain analyses lies in their ability to refer to the 'same' location across studies, below the level of sulci and gyri (e.g., focusing on the portion of the inferior frontal gyrus that corresponds to BA 45). However, because, as discussed above, the borders of cytoarchitectonic zones cannot be determined in any given individual brain, these probabilistic maps run into some of the same limitations as other group-based methods. (If it were possible to define cytoarchitectonic areas in individual brains, this would likely lead to tremendous progress in understanding the neural substrates of cognitive processes in humans.)

Probabilistic cytoarchitectonic maps can be used in at least two ways. First, they can be used as anatomical ROIs from which BOLD (blood-oxygen-level-dependent) signal is extracted (as in Figure 2). And second, for any given cluster that emerges in a traditional random-effects analysis, the probability of that cluster falling within a particular cytoarchitectonic area can be calculated. Each of these analysis methods is problematic. With respect to using these maps as aROIs, the following problem arises: due to high variability in the precise locations of these areas across the 10 post-mortem brains that were used to create these maps, probabilistic maps for adjacent areas are highly overlapping. For instance, considering the two BAs that constitute Broca's area, BA 44 and BA 45: 43.0% of BA 44 (the opercular part of the IFG) voxels also belong to BA 45 (the triangular part of the IFG), and 51.3% of BA 45 voxels also belong to BA 44. As a result, in order to separate the response from adjacent BAs, the corresponding probabilistic maps need to be thresholded to reduce/minimize overlap. For example, one may consider only those voxels that correspond to the target BA in most of the 10 brains. However, such thresholding procedures result in a loss of information on the extent of variability in the locations

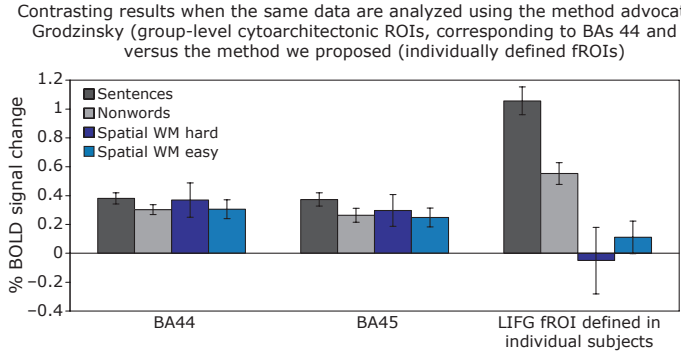

Fig 2. Responses of Brodmann Areas 44 and 45 (defined with the probabilistic cytoarchitectonic maps[12] advocated by Grodzinsky; Amunts et al. 1999) and a functional ROI in the left inferior frontal gyrus, defined in individual subjects, as proposed by F&K. Responses are to the language localizer task and a spatial working memory task (remembering more vs. fewer spatial locations on a grid). For the response of the LIFG fROI to the localizer conditions, the individual subjects' ROIs are defined using all but the first functional run, and the responses are estimated using the first run's data to avoid non-independence issues (e.g., Kriegeskorte et al. 2009).

of cytoarchitectonic areas across subjects, which is what these maps are trying to capture in the first place (and what functional ROIs are designed to overcome). In addition, using the same set of voxels in stereotaxic space (corresponding to the BA map, thresholded or not) will underestimate functional selectivity (see Fig. 2 and Fedorenko et al. forthcoming, for empirical support).

With respect to calculating the probability of an activation cluster falling within a particular cytoarchitectonic area (e.g., the probability of the 'syntactic movement' cluster in Santi and Grodzinsky 2007, falling within BA45), the problem arises when interpreting the results. It seems too arbitrary to set a hard threshold for counting an activation cluster as falling within the borders of some BA (e.g., a cluster falls within the borders of a BA if the probability value is higher than 50%), or for a minimal difference in the probabilities of the cluster falling within a particular BA and not any of the adjacent BAs (e.g., the probability of the cluster falling within BA45 is at least 25% higher than that of the cluster falling within BA44 and/or other adjacent BAs). One rigorous way to determine whether activations for a particular contrast reflect activity of a particular BA would involve examining the topography of individual activations falling around the area of interest and comparing it to the spatial distribution to that of the cytoarchitectonic area in question in the 10 brains that were used in creating the maps, to see whether these activations are likely to be generated by the same underlying source. One problem with this approach is that given the extent of the variability, the sample of only 10 brains may be too small to reliably estimate the topography of the distribution in the locations of cytoarchitectonic areas. Other possibilities are worth considering, but regardless of what particular procedure is chosen for determining whether some activation cluster does or does not fall within the boundaries of a particular BA, this procedure needs to be standardized so that the researcher is constrained in interpreting the data (see Fedorenko et al. forthcoming, for additional discussion).

Grodzinsky also argues that the reason why probabilistic cytoarchitectonic maps are not used in studying high-level visual areas (like the fusiform face area, Kanwisher et al. 1997) is that no such maps exist for the fusiform gyrus. However, this is misguided. For the reasons discussed above and given how successful the application of functional localizers has been in studying high level vision, researchers studying vision should not sacrifice the ability to precisely and reliably identify the relevant parts of cortex in each individual brain for a group-level probabilistic representation of the underlying cytoarchitecture.

A small aside related to this point is worth making. On p. 615, Grodzinsky points out the following, 'Brodmann's Area 21, which contains the fusiform, has not yet lent itself to cytoarchitectonic parsing, remaining one of the largest cytoarchitectonically uniform chunks in the human brain'. And then a few lines down, 'The fusiform gyrus constitutes a single, large, cytoarchitectonic area'. The presence of well characterized subregions within the fusiform gyrus (e.g., the fusiform face area, the fusiform body area) that are functionally distinct raises a general concern with treating borders of cytoarchitectonic areas (e.g., based on Brodmann's parcellation, adopted by Zilles, Amunts, and colleagues) as borders of functional modules, which is what Grodzinsky advocates (see e.g., section 5 of his response). In other words, borders of BAs 44 and 45 *may* indicate boundaries between internally homogeneous functionally distinct regions. However, it is also possible that subregions within these areas exist which have distinct functional properties (see Amunts et al. 1999, 2010, for a discussion of some evidence for such subregions based on cyto-/myelo-/receptor-architecture), just like they exist within the cytoarchitectonically uniform BA 21. Consequently, the concern that Grodzinsky raises with respect to treating functionally defined ROIs as functionally uniform applies in a very similar way to anatomical ROIs

based on the borders of cytoarchitectonic zones. Thus, by pointing out that a region known to be functionally heterogeneous – the fusiform gyrus – constitutes a region that has not yet been cytoarchitectonically subdivided, Grodzinsky indirectly reveals a key limitation with restricting oneself to only cytoarchitectonically defined ROIs.

Finally, the relationship between microarchitecture (e.g., cytoarchitecture, myeloarchitecture, differential connectivity across cortical regions) and functional activations is a crucial unanswered question. In some domains (like vision), this question can be addressed by relating fMRI activations in animals to the cytoarchitectonic properties of the relevant cortical regions in their brains, post-mortem. With respect to language (and other human-specific cognitive processes) this question is more difficult to address. We hope that, with an increasingly widespread use of fMRI and the development of brain donation programs, this will some day become possible. It is worth noting that paradigms that elicit robust functional activations at the individual subject level will be of particular use in this enterprise.

3.4. How probabilistic cytoarchitectonic maps can be combined with the individual-subjects fROI approach

In Fedorenko et al. (forthcoming), we suggested one way in which probabilistic cytoarchitectonic maps can be combined with the individual-subject fROI approach. In particular, one challenge for individual-subjects fROI methods is establishing correspondence between activations in one brain and another brain (i.e., deciding what part of the activation constitutes a 'region' and how these regions correspond across subjects). In Fedorenko et al. (2010), we proposed a novel solution to this problem: a Group-constrained Subject-Specific (GSS) analysis method. In this method, thresholded individual maps are overlaid in the common space, and the topography of the resulting probabilistic overlap map is used to derive a set of 'group-level functional partitions'. These partitions are then intersected with each individual activation map to define subject-specific fROIs (see J. Julian, E. Fedorenko, N. Kanwisher, unpublished data, for a demonstration of the efficacy of this method for well-characterized visual regions, like the fusiform face area). However, an alternative way to constrain the selection of subject-specific voxels for a particular region would be to use anatomical partitions, like, for example, the probabilistic cytoarchitectonic maps. The critical aspect of this analysis is the intersection of, e.g., a map for BA45 with individual activation maps rather than using the same set of voxels across subjects (see Fedorenko et al. forthcoming, for an empirical demonstration).

4. Conclusions

In this response, we tried to clarify our position on the current state of affairs in neurolinguistic research, as well as point out some problems with the way Grodzinsky paints this picture. We then summarized the space of possibilities with respect to fMRI analyses, addressed the concern that Grodzinsky raised with respect to the use of functional localizers (referring the reader to additional papers where these and other issues are discussed at length), and outlined some concerns with the use of probabilistic cytoarchitectonic maps.

We make two additional general points in concluding.

1 Researchers working on questions of the brain basis of language may vary in how satisfied they are with the progress that the field has made over the last two or three decades. From the point of view of many researchers in the fields of

psycholinguistics/sentence processing, there is a large gap between our understanding of the cognitive architecture of language (based on what we have learned from behavioral and computational modeling work) vs. our understanding of the neural implementation of linguistic processes. Many important questions about the brain basis of language have not yet been answered. As a result, new methods should be welcome, especially if they can be shown to outperform existing methods. In addition to the higher sensitivity and selectivity that the individual-subjects functional localization approach affords, this method has another key advantage. In particular, it enables us to refer to the 'same' region(s) consistently across studies and across labs, thus making possible a cumulative research enterprise where functional profiles of a set of regions are characterized in detail on the basis of many studies across researchers and labs. Only such detailed and systematic characterization may eventually enable us to infer the computations these regions perform, or the representations they store and manipulate.

2 One interesting point Grodzinsky raises is the need to be guided by a particular theoretical framework in investigating the brain basis of language. He says on p. 613, '… FK are silent on the model that guides their exploration'. Although having a particular 'model' of a cognitive process may be useful in thinking about the neural basis of the cognitive process of interest (as long as alternative models are entertained), with respect to language it is unclear what 'models' Grodzinsky has in mind. In our opinion, the best strategy for making progress in understanding the brain basis of language is to be aware of the work that has taken place in psycholinguistics over the last fifty years. Across many subfields of psycholinguistics, important advances have been made and some quite detailed proposals have been articulated concerning various linguistic processes (e.g., syntactic complexity – Gibson 1998; Lewis et al. 2006; Levy 2008; or lexical access – Levelt 1989; Caramazza 1997). If we are to make progress in understanding the neural basis of language, and in establishing a stronger connection between the behavioral and brain-based work, it will be important to be broadly aware of the factors that have been shown to affect language comprehension and production processes.

We have focused here on the theoretical arguments about why individual-subjects analyses may lead to a clearer picture of the functional specificity of language regions in the brain. But of course the proof is in the pudding. So, as a small taste of things to come, we end by showing some pilot data that uses our method to test domain-specificity of a language-sensitive region located in the LIFG, and we contrast these results with analyses of the same data using the group-level cytoarchitectonic ROIs advocated by Grodzinsky. As revealed in Figure 2, the sensitivity and specificity of the results obtained from our method are vastly stronger than those obtained from the cytoarchitectonic ROI-based method. If Grodzinsky and others prefer to obtain results like those shown in the first two sets of bars in Figure 2 they are free to do so. We find results like those shown in the last set of bars more informative. Indeed, it is a crowning irony of Grodzinsky's critique that the very method we advocate has the highest chance of detecting functional specificity of the kind he hypothesizes if such specificity is present in the language system.

All of the tools we developed are available on the web (http://web.mit.edu/evelina9/www/funcloc). We invite Grodzinsky, and the rest of the scientific community, to give them a try.

Short Biographies

Evelina Fedorenko is a postdoctoral associate in the Kanwisher Lab at the McGovern Institute for Brain Research at MIT. She received her graduate training at MIT with Ted Gibson and Nancy Kanwisher. Her current work is aimed at investigating the neural basis of language using fMRI and other techniques. In particular, she is applying the functional localization approach, which has been successful in other domains, to language. This approach involves functionally defining a set of language-sensitive regions in individual brains, and subsequently investigating the functional profiles of these regions. Functional localization in individual subjects circumvents problems of inter-subject anatomical variability by allowing to pool data from corresponding functional regions across subjects, and enables more efficient knowledge accumulation by allowing researchers to consistently refer to the same regions across studies and labs. In addition to the questions of the neural basis of language, Fedorenko is interested in the functional architecture of the language system more generally, working on a variety of questions related to the nature of language comprehension and production mechanisms.

Nancy Kanwisher is the Ellen Swallow Richard Professor in the Department of Brain & Cognitive Sciences at MIT, and Investigator at MIT's McGovern Institute for Brain Research. She received her B.S. in 1980 and her PhD in 1986, both from MIT. Kanwisher taught at UCLA and Harvard before returning to MIT in 1997. Kanwisher's lab has contributed to the identification and characterization of a number of regions in the human brain that conduct very specific cognitive functions: four are involved in the visual perception of specific kinds of stimuli (faces, places, bodies, and words), and another is selectively engaged in inferring the contents of another person's thoughts. Kanwisher received a Troland Research Award from the National Academy of Sciences in 1999, a MacVicar Faculty Fellow teaching Award from MIT in 2002, and the Golden Brain Award from the Minerva Foundation in 2007. She was elected as a member of the National Academy of Sciences in 2005 and the American Academy of Arts and Sciences in 2009.

Acknowledgement

We would like to thank Ted Gibson for helpful comments on this paper.

Notes

* Correspondence address: Evelina Fedorenko, McGovern Institute for Brain Research, MIT, 43 Vassar Street, 46-4141C, Cambridge, MA 02139, USA. E-mail: evelina9@mit.edu

[1] Technically, these two issues could be non-independent if one were to argue that the need for a new method is obsolete because we, as a field, have already arrived at a complete understanding of the neural basis of language using the currently available methods. Although researchers may vary in how satisfied they are with the progress that has been made in understanding the brain basis of language in the last several decades, it seems unlikely that anyone could argue that all the key questions have been answered. As a result, we discuss the two issues separately.

[2] All the arguments laid out in this section apply in a similar way to patient studies.

[3] Note that the contrasts Grodzinsky discusses differ not only with respect to syntactic movement (as acknowledged by Grodzinsky in passing on p. 611). See also Section 2.2.1.3 below for additional discussion.

[4] Note also that some of these contrasts are not matched for difficulty. For example, consider the contrast between subject- and object-extracted relative clauses, on the one hand, and subject- and object-modifying relative clauses on the other hand. Whereas object-extracted relative clauses (II and IV in table 2 in Grodzinsky 2010) cause processing difficulty relative to subject-extracted relative clauses (I and III in table 2), no such difference exists between subject-modifying relative clauses (I and II in table 2) and object-modifying relative clauses (III and IV in table 2;

see e.g., Gibson and Thomas 1997, for evidence from similar constructions). In order to meaningfully interpret the presence of an effect for one manipulation in the absence of an effect for another manipulation, it is important to match the relevant manipulations for difficulty.

[5] It is also worth clarifying a misunderstanding apparent in Grodzinsky's response. On p. 616, Grodzinsky says the following, 'Everyone seems to follow a research program that seeks structure by delineating the functions of a brain region and exploring each independently. We assume that the functions of each region are independent, until a theory that unifies them is proposed. This approach, which allows regional multifunctional modularity, is in line with the one that Saxe et al. (2006) recommend: study functions separately from a coherent theoretical perspective, hope that higher-resolution devices and maps are in the making, but first pursue each separately, and then try to understand how they relate to one another'. This is not what Saxe et al. recommend. Instead they recommend studying *brain regions* separately, carefully characterizing each region's functional profile by examining the region's response to a range of stimuli/tasks. They do not advocate focusing on a subset of activations observed in a region and not paying attention to activations that do not fit with a particular theoretical hypothesis about the region's function(s).

[6] Nicol et al. (1994) responded to McKoon and colleagues. However, the objections they raise (with respect to some methodological choices made by McKoon and colleagues) do not undermine the conclusions that McKoon and colleagues draw based on their experiments: namely, that the 'reactivation' effects were due to an experimental confound.

[7] This space of possibilities (with respect to the units of analysis in fMRI) is orthogonal to another dimension in fMRI analyses. In particular, with respect to all the different analyses laid out in Figure 1, it is possible to use a host of different designs, such as a simple subtraction (e.g., comparing the brain's response to reading sentences vs. reading strings of unconnected words), parametric manipulations (e.g., examining the brain's response to reading words that vary along some continuous dimension, like lexical frequency), multivoxel pattern analyses (e.g., examining the brain's response to one type of linguistic stimulus vs. another to see whether the distinction is represented in the pattern of activity across multiple voxels), neural adaptation (e.g., examining the brains response to stimuli that are presented in succession and share some relevant property to see whether the response to the second stimulus is decreased), etc.

[8] These papers typically use *ad hoc* contrasts, not independently validated localizer contrasts.

[9] Most of the questions/issues Grodzinsky raises at the beginning of the relevant section (section 8) have been addressed in Fedorenko et al. (2010).

[10] Although Grodzinsky criticizes the localizer contrast we chose, he actually used a similar contrast in his own previous work to functionally define ROIs in individual subjects (Ben-Shachar et al. 2004).

[11] The novel analysis we developed for performing group-level analyses, taking individual subjects' activations into consideration (GSS analysis; this method was originally introduced with the abbreviation GcSS) will be perfectly suited for developing new localizers based on additional functional contrasts, as it is more sensitive than the traditional random-effects group analysis (E. Fedorenko and N. Kanwisher, unpublished data, J. Julian, E. Fedorenko, N. Kanwisher, unpublished data).

[12] For the analysis presented in Figure 2, we used unthresholded probabilistic cytoarchitectonic maps (i.e., if a voxel corresponded to the relevant cytoarchitectonic area in *at least one* of the 10 brains, it was included). In Fedorenko et al. (forthcoming), we show that the results are similar for probabilistic maps thresholded at different (even very conservative) levels.

Works Cited

Amunts, K., O. Kedo, M. Kindler, P. Pieperhoff, H. Mohlberg, N. J. Shah, U. Habel, F. Schneider and K. Zilles. 2005. Cytoarchitectonic mapping of the human amygdala, hippocampal region and entorhinal cortex: intersubject variability and probability maps. Anatomy and Embryology 210(5–6). 343–52.

——, M. Lenzen, A. Friederici, A. Schleicher, P. Morosan, N. Palomero-Gallagher, and K. Zilles. 2010. Broca's region: novel organizational principles and multiple receptor mapping. PLoS Biology 8(9). 1–16.

——, A. Malikovic, H. Mohlberg, T. Schormann, and K. Zilles. 2000. Brodmann's areas 17 and 18 brought into stereotaxic space-where and how variable? Neuroimage 11(1). 66–84.

——, A. Schleicher, U. Burgel, H. Mohlberg, H. B. Uylings, and K. Zilles. 1999. Broca's region revisited: cytoarchitecture and intersubject variabilit. Journal of Comparative Neurology 412(2). 319–41.

Ben-Shachar, M., D. Palti, and Y. Grodzinsky. 2004. Neural correlates of syntactic movement: converging evidence from two fMRI experiments. Neuroimage 21(4). 1320–36.

Blumstein, S. E., E. B. Myers, and J. Rissman. 2005. The perception of voice onset time: an fMRI investigation of phonetic category structure. Journal of Cognitive Neuroscience 17(9). 1353–66.

Bresnan, J. 1982. The mental representation of grammatical relations. Cambridge, MA: MIT Press.

Caramazza, A. 1997. How many levels of processing are there in lexical access? Cognitive Neuropsychology 14. 177–208.

Caspers, S., S. Geyer, A. Schleicher, H. Mohlberg, K. Amunts, and K. Zilles. 2006. The human inferior parietal cortex: cytoarchitectonic parcellation and interindividual variability. Neuroimage 33(2). 430–48.
Chomsky, N. 1965. Aspects of the theory of syntax. Cambridge: The MIT Press.
——. 1973. Conditions on transformations. A festschrift for Morris Halle, ed. by S. Anderson and P. Kiparsky, 232–86. New York: Holt, Rinehart & Winston.
——. 1981. Lectures on government and binding: the Pisa lectures. Holland: Foris Publications.
——. 1993. A minimalist program for linguistic theory. The view from Building 20, ed. by K. Hale and S. J. Keyser, 1–52. Cambridge, MA: MIT Press.
Clark, V. P., E. Courchesne, and M. Grafe. 1992. *In vivo* myeloarchitectonic analysis of human striate and extrastriate cortex using magnetic resonance imaging. Cerebral Cortex 2. 417–24.
Dehaene, S., E. Spelke, P. Pinel, R. Stanescu, and S. Tsivkin. 1999. Sources of mathematical thinking: behavioral and brain-imaging evidence. Science 284(5416). 970–4.
Duncan, J. 2001. An adaptive coding model of neural function in prefrontal cortex. Nature Reviews Neuroscience 2(11). 820–9.
Fedorenko, E., E. Gibson, and D. Rohde. 2006. The nature of working memory capacity in sentence comprehension: evidence against domain-specific working memory resources. Journal of Memory and Language 54(4). 541–53.
——, ——, and ——. 2007. The nature of working memory in linguistic, arithmetic and spatial integration processes. Journal of Memory and Language 56(2). 246–69.
——, P.-J. Hsieh, A. Nieto-Castañon, S. Whitfield-Gabrieli, and N. Kanwisher. 2010. A new method for fMRI investigations of language: defining ROIs functionally in individual subjects. Journal of Neurophysiology 104. 1177–94.
——, and K. Kanwisher. 2009. Neuroimaging of language: why hasn't a clearer picture emerged? Language & Linguistics Compass 3(4). 839–65.
——, A. Nieto-Castanon, and K. Kanwisher. Forthcoming. Syntactic processing in the human brain: what we know, what we don't know, and a suggestion for how to proceed. Brain & Language.
Friston, K. J., P. Rotshtein, J. J. Geng, P. Sterzer, and R. N. Henson. 2006. A critique of functional localizers. Neuroimage 30(4). 1077–87.
Gennari, S. P., and M. C. MacDonald. 2009. Linking production and comprehension processes: the case of relative clauses. Cognition 111. 1–23.
Gibson, E. 1998. Linguistic complexity: locality of syntactic dependencies. Cognition 68. 1–76.
——, and J. Thomas. 1997. The complexity of nested structures in English: evidence for the syntactic prediction locality theory of linguistic complexity. Massachusetts Institute of Technology, Unpublished Manuscript.
Gordon, P. C., R. Hendrik, and W. H. Levine. 2002. Memory-load interference in syntactic processing. Psychological Science 13(5). 425–30.
Grodzinsky, Y. 2010. The picture of the linguistic brain: how sharp can it be? Reply to Fedorenko & Kanwisher Language and Linguistics Compass 4(8). 605–22.
——, and A. Santi. 2008. The battle for Broca's region. Trends in Cognitive Sciences 12(12). 474–80.
Hagoort, P., G. Baggio, and R. M. Willems. 2009. Semantic unification. The cognitive neurosciences IV, ed. by M. Gazzaniga, 819–36. Cambridge, MA: MIT Press.
——, L. Hald, M. Bastiaansen, and K. M. Petersson. 2004. Integration of word meaning and world knowledge in language comprehension. Science 304(5669). 438–41.
Hickok, G., K. Okada, and J. T. Serences. 2009. Area Spt in the human planum temporale supports sensory-motor integration for speech processing. Journal of Neurophysiology 101. 2725–32.
Iwamura, Y., M. Tanaka, M. Sakamoto, and O. Hikosaka. 1983. Converging patterns of finger representation and complex response properties of neurons in area 1 of the first somatosensory cortex of the conscious monkey. Experimental Brain Research 51(3). 327–37.
January, D., J. C. Trueswell, and S. L. Thompson-Schill. 2009. Co-localization of stroop and syntactic ambiguity resolution in Broca's area: implications for the neural basis of sentence processing. Journal of Cognitive Neuroscience 21(12). 2434–44.
Kanwisher, N. 2010. Functional specificity in the human brain: a window into the functional architecture of the mind. Proceedings of the National Academy of Sciences 107(25). 11163–70.
——, J. McDermott, and M. M. Chun. 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. Journal of Neuroscience 17(11). 4302–11.
Kriegeskorte, N., W. K. Simmons, P. S. Bellgowan, and C. I. Baker. 2009. Circular analysis in systems neuroscience: the dangers of double dipping. Nature Neuroscience 12(5). 535–40.
Kuperberg, G. R., P. J. Holcomb, T. Sitnikova, D. Greve, A. M. Dale, and D. Caplan. 2003. Distinct patterns of neural modulation during the processing of conceptual and syntactic anomalies. Journal of Cognitive Neuroscience 15(2). 272–93.
Levelt, W. J. M. 1989. Speaking. Cambridge, MA: The MIT Press.

Levitin, D. J., and V. Menon. 2003. Musical structure is processed in ''language'' areas of the brain: a possible role for Brodmann area 47 in temporal coherence. Neuroimage 20(4). 2142–52.
Levy, R. 2008. Expectation-based syntactic comprehension. Cognition 106(3). 1126–77.
Lewis, R. L., S. Vasishth, and J. A. Van Dyke. 2006. Computational principles of working memory in sentence comprehension. Trends in Cognitive Sciences 10. 44–54.
Lindenberg, R., H. Fangerau, and R. J. Seitz. 2007. Broca's area'' as a collective term? Brain and Language 102(1). 22–9.
Maess, B., S. Koelsch, T. C. Gunter, and A. D. Friederici. 2001. Musical syntax is processed in Broca's area: an MEG study. Nature Neuroscience 4. 540–5.
Matelli, M., G. Luppino, and G. Rizzolatti. 1991. Architecture of superior and mesial area 6 and the adjacent cingulate cortex in the macaque monkey. Journal of Comparative Neurology 311(4). 445–62.
McElree, B., S. Foraker, and L. Dyer. 2003. Memory structures that subserve sentence comprehension. Journal of Memory and Language 48. 67–91.
McKoon, G., D. W. Allbritton, and R. Ratcliff. 1996. Sentential context effects on lexical decisions with a cross-modal instead of an all-visual procedure. Journal of Experimental Psychology: Learning, Memory, and Cognition 22. 1494–7.
——, and R. Ratcliff. 1994. Sentential context and on-line lexical decision. Journal of Experimental Psychology: Learning, Memory, and Cognition 20. 1239–43.
Morosan, P., J. Rademacher, A. Schleicher, K. Amunts, T. Schormann, and K. Zilles. 2001. Human primary auditory cortex: cytoarchitectonic subdivisions and mapping into a spatial reference system. Neuroimage 13(4). 684–701.
Myers, E. B., S. E. Blumstein, E. Walsh, and J. Eliassen. 2009. Inferior frontal regions underlie the perception of phonetic category invariance. Psychological Science 20(7). 895–903.
Nickerson, R. S. 1998. Confirmation bias: a ubiquitous phenomenon in many guises. Review of General Psychology 2. 175–220.
Nicol, J. L., J. D. Fodor, and D. Swinney. 1994. Using cross-modal lexical decision tasks to investigate sentence processing. Journal of Experimental Psychology: Learning, Memory, and Cognition 20. 1229–38.
——, and D. Swinney. 1989. The role of structure in coreference assignment during sentence comprehension. Journal of Psycholinguistic Research 18. 5–20.
Nieto-Castañon, A., E. Fedorenko, and N. Kanwisher. Forthcoming. Functional localizers increase sensitivity and selectivity.
——, S. S. Ghosh, J. A. Tourville, and F. H. Guenther. 2003. Region of interest based analysis of functional imaging data. Neuroimage 19(4). 1303–16.
Ojemann, G. A. 1991. Cortical organization of language. The Journal of Neuroscience 11(8). 2281–7.
Owen, A. M., K. M. McMillan, A. R. Laird, and E. Bullmore. 2005. N-Back working memory paradigm: a meta-analysis of normative functional neuroimaging studies. Human Brain Mapping 25. 46–59.
Pollard, C., and I. A. Sag. 1994. Head-driven phrase structure grammar. Chicago: University of Chicago Press.
Rodd, J. M., M. H. Davis, and I. S. Johnsrude. 2005. The neural mechanisms of speech comprehension: fMRI studies of semantic ambiguity. Cerebral Cortex 15(8). 1261–9.
Rozzi, S., P. F. Ferrari, L. Bonini, G. Rizzolatti, and L. Fogassi. 2008. Functional organization of inferior parietal lobule convexity in the macaque monkey: electrophysiological characterization of motor, sensory and mirror responses and their correlation with cytoarchitectonic areas. European Journal of Neuroscience 28(8). 1569–88.
Santi, A., and Y. Grodzinsky. 2007. Working memory and syntax interact in Broca's area. Neuroimage 37(1). 8–17.
Saxe, R., M. Brett, and N. Kanwisher. 2006. Divide and conquer: a defense of functional localizers. Neuroimage 30(4). 1088–96; discussion 1097–9.
——, and ——. 2006. It's the thought that counts: specific brain regions for one component of theory of mind. Psychological Science 17(8). 692–9.
Schnur, T. T., M. F. Schwartz, D. Y. Kimberg, E. Hirshorn, H. B. Coslett, and S. L. Thompson-Schill. 2009. Localizing interference during naming: convergent neuroimaging and neuropsychological evidence for the function of Broca's area. Proceedings of the National Academy of Sciences of the United States of America 106(1). 322–7.
Vigneau, M., V. Beaucousin, P. Y. Hervé, H. Duffau, F. Crivello, O. Houdé, B. Mazoyer, and N. Tzourio-Mazoyer. 2006. Meta-analyzing left hemisphere language areas: phonology, semantics, and sentence processing. Neuroimage 30(4). 1414–32.
Wason, P. C. 1960. On the failure to eliminate hypotheses in a conceptual task. Quarterly Journal of Experimental Psychology 12. 129–40.
Willems, R. M., and P. Hagoort. 2009. Broca's region: battles are not won by ignoring half of the facts. Trends in Cognitive Science 13(3). 101; author reply 102.