This document contains the exchange that took place between Greg Hickok and Ev Fedorenko & Nancy Kanwisher on the TalkingBrains blog (http://www.talkingbrains.org) in Aug-Sept 2010.

There are six parts to this exchange:

1) Greg Hickok – August 24, 2010
2) Ev Fedorenko & Nancy Kanwisher – August 31, 2010
3) Greg Hickok – August 31, 2010
4) Ev Fedorenko & Nancy Kanwisher – August 31, 2010
5) Greg Hickok – September 1, 2010
6) Ev Fedorenko & Nancy Kanwisher – September 1, 2010

On the blog, Part 1 can be found at:
http://www.talkingbrains.org/2010/08/neuroimaging-of-language-why-hasnt.html

Parts 2-6 can be found at:
http://www.talkingbrains.org/2010/08/response-from-fedorenko-kanwisher.html

PART 1: August 24, 2010

Greg Hickok wrote:

Neuroimaging of language: Why hasn't a clearer picture emerged?

This is the question raised in a paper by Evelina Fedorenko and Nancy Kanwisher published last year in *Language and Linguistics Compass*. The main point that they want to make is that language neuroimagers need to stop doing group studies and start doing functional localization in individual subjects, like the vision folks do. I don't disagree at all; e.g., see http://www.talkingbrains.org/2007/07/problem-with-group-studies-in-fmri.html. In fact, we have used individual subject analyses in several of our papers (e.g., Okada & Hickok, 2006; Okada et al., in press).

What I found a bit on the irritating side though was the extremely dim and distressingly myopic view of progress in the field of the neural basis of language. They start by stating that two questions have "driven dozens of studies on the neural basis of language published in the last several decades: (i) Are distinct cortical regions engaged in different aspects of language? (ii) Are regions engaged in language processing specific to the domain of language?"

And they suggest that "Neuroimaging has not yet provided clear answers to either question".

Regarding question one, there is strong evidence from functional imaging regarding the involvement of distinct cortical regions/circuits in phonemic (STS), lexical-semantic (MTG), prosodic (anterior dorsal STG), and higher-level combinatorial processes (anterior temporal/inferior frontal regions). Additional circuits have been delineated that support auditory-motor integration and auditory/phonological short-term memory. Here are some relevant reviews of this literature: Binder et al., 2000; Hickok and Poeppel, 2007; Indefrey and Levelt, 2004.

Regarding question two, several studies have clearly identified voice-specific responses in the STG (Belin et al. 2000), higher-level speech specific responses in the STS (Scott et al. 2000; Okada et al. 2010), and even what we might as well call the anterior temporal lobe sentence area, given how selective it is to the perception of sentence-level stimulation (Humphries et al., 2006; Humphries et al., 2005; Humphries et al., 2001; Rogalsky and Hickok, 2009; Vandenberghe et al., 2002). But more importantly, some of us have moved beyond the specificity issue with the aim of trying to identify the circuits and computations involved in a given process whether or not it is special to speech.

So F&K are a bit misinformed in regarding the contribution of neuroimaging to the questions they raise.

Even worse, though, is their summary of "where things stand" concerning our understanding of the "neural basis of language" -- a rather sweeping domain, especially if they have in mind only the two questions they raise at the outset. Nonetheless, concerning the Neural Basis of Language they emphasize:

1. that the 19th century idea that Broca's area = speech production and Wernicke's area = speech comprehension doesn't hold up to modern data

2. that left frontal regions activate to a variety of language tasks, and indeed even non-linguistic tasks

3. that regions outside of the traditional peri-Sylvian cortex activate during language processing

4. that meta analyses show lots of overlap between language tasks

This is frankly a pathetic summary of the state of the field and a pretentious starting point for the methodological schooling that F&K provide in the following sections of their paper.

Completely ignored in this summary is (i) a body of work showing that much of the confusion (and overlapping activations) evaporates if one is careful about task selection (Hickok & Poeppel, 2007), (ii) convergence on the involvement of the STS in phonemic level processes in speech perception (Leibenthal et al., 2005; Scott & Johnsrude, 2003; Hickok & Poeppel, 2007), (iii) convergence on the idea of a dual stream architecture in language system (Hickok & Poeppel, 2007; Rauchecker & Scott, 2009), (iv) recent progress in mapping the circuit that supports sensory-motor integration in speech processing (Golfinopoulus et al. 2009; Hickok et al., 2009), (v) progress in understanding the basis of hemispheric asymmetries for acoustic and phonemic processing in auditory cortex (Boemio, et al. 2005); Zatorre, et al. 2002) , (vi) convergence on the idea that anterior temporal regions support some aspect of sentence-level processing (the linguistic equivalent of the FFA), (vii) convergence on the relation between sensory-motor circuits and phonological short-term memory (Buchsbaum et al. 2008; Postle, 2006)... I could go on.

Yes, there is still plenty of murkiness, much of it surrounding the function of Broca's area, and yes, individual subject analyses would be helpful, but it is not a magic bullet (e.g., task selection is more important in my view) -- e.g., I'm willing to bet that the vision folks still have some work to do -- and the existence of murkiness doesn't justify the characterization of an entire field as failing to make progress due to methodological ineptness. This kind of argumentation was prominent in another of Fedorenko's papers that was featured prominently on this blog (http://www.talkingbrains.org/2010/06/egregious-act-of-methodological.html). It's a bit disturbing to see it showing up again.

F&K's paper has generated a more formal (i.e., published) response by Grodzinsky (2010) who is critical of their take on the field as well but for different reasons. Definitely worth a look.

The field of the neural basis of language has made significant progress in the last several years, despite what F&K assert.

References

Fedorenko, E., & Kanwisher, N. (2009). Neuroimaging of Language: Why Hasn't a Clearer Picture Emerged? *Language and Linguistics Compass, 3* (4), 839-865.

Binder, J.R., Frost, J.A., Hammeke, T.A., Bellgowan, P.S., Springer, J.A., Kaufman, J.N., and Possing, E.T. (2000). Human temporal lobe activation by speech and nonspeech sounds. Cerebral Cortex 10, 512-528.

Belin, P., Zatorre, R.J., Lafaille, P., Ahad, P., and Pike, B. (2000). Voice-selective areas in human auditory cortex. Nature 403, 309-312.

Buchsbaum, B.R., and D'Esposito, M. (2008). The search for the phonological store: from loop to convolution. J Cogn Neurosci 20, 762-778.
Golfinopoulos, E., Tourville, J.A., and Guenther, F.H. (2009). The integration of large-scale neural network modeling and functional brain imaging in speech motor control. Neuroimage.
Grodzinsky, Yosef. 2010. A clearer view of the linguistic brain: reply to Fedorenko and Kanwisher. Language and Linguistic Compass, 4, pp. 605-622.
Hickok, G., and Poeppel, D. (2007). The cortical organization of speech processing. Nat Rev Neurosci 8, 393-402.
Humphries, C., Binder, J.R., Medler, D.A., and Liebenthal, E. (2006). Syntactic and semantic modulation of neural activity during auditory sentence comprehension. J Cogn Neurosci 18, 665-679.
Humphries, C., Love, T., Swinney, D., and Hickok, G. (2005). Response of anterior temporal cortex to syntactic and prosodic manipulations during sentence processing. Human Brain Mapping 26, 128-138.
Humphries, C., Willard, K., Buchsbaum, B., and Hickok, G. (2001). Role of anterior temporal cortex in auditory sentence comprehension: An fMRI study. Neuroreport 12, 1749-1752.
Indefrey, P., and Levelt, W.J. (2004). The spatial and temporal signatures of word production components. Cognition 92, 101-144.
Liebenthal, E., Binder, J.R., Spitzer, S.M., Possing, E.T., and Medler, D.A. (2005). Neural substrates of phonemic perception. Cereb Cortex 15, 1621-1631.
Okada, K., and Hickok, G. (2006). Identification of lexical-phonological networks in the superior temporal sulcus using fMRI. Neuroreport 17, 1293-1296.
Okada, K., Rong, F., Venezia, J., Matchin, W., Hsieh, I.H., Saberi, K., Serences, J.T., and Hickok, G. (in press). Hierarchical Organization of Human Auditory Cortex: Evidence from Acoustic Invariance in the Response to Intelligible Speech. Cereb Cortex.
Postle, B.R. (2006). Working memory as an emergent property of the mind and brain. Neuroscience 139, 23-38.
Rauschecker, J.P., and Scott, S.K. (2009). Maps and streams in the auditory cortex: nonhuman primates illuminate human speech processing. Nat Neurosci 12, 718-724.
Rogalsky, C., and Hickok, G. (2009). Selective Attention to Semantic and Syntactic Features Modulates Sentence Processing Networks in Anterior Temporal Cortex. Cereb Cortex 19, 786-796.
Scott, S.K., and Johnsrude, I.S. (2003). The neuroanatomical and functional organization of speech perception. Trends Neurosci 26, 100-107.
Scott, S.K., Blank, C.C., Rosen, S., and Wise, R.J.S. (2000). Identification of a pathway for intelligible speech in the left temporal lobe. Brain 123, 2400-2406.
Vandenberghe, R., Nobre, A.C., and Price, C.J. (2002). The response of left temporal cortex to sentences. J Cogn Neurosci 14, 550-560.

David Poeppel commented:

I am a grumpy middle-aged guy, and I tend to be pretty critical of basically everything I come across -- but even I am not as negative about the state of the field as Fedorenko and Kanwisher ... I think they seriously underestimate the sophistication of some of the current work, and Greg has already pointed out that there is a literature dealing with their specific complaints.

Let me amplify one point. While functional imaging -- and as a consequence functional localization -- plays a dominant role in the field, a lot of research effort is going into increasingly detailed analyses of the tasks, and specifically the computational subroutines underlying the execution of the experimental tasks. We are presumably all pulling in the same direction, aiming for theoretically well motivated, computationally explicit, and neurobiologically solidly grounded accounts of speech and language, so I would welcome new approaches to understand things more deeply -- but the suggestions articulated by F&K are not yet mature enough to motivate me to adopt that perspective, given the tangible progress we have already made.

PART 2: August 31, 2010

Ev Fedorenko and Nancy Kanwisher wrote:

Dear Greg and David:

Thank you for your comments on our 2009 paper (Fedorenko & Kanwisher, 2009, *Language and Linguistics Compass*). As you stated in your post on August 24, 2010, and as became even clearer from the email exchange that took place the following day, you agree with Nancy Kanwisher and me on the importance of taking into consideration anatomical and functional variability in fMRI research in any domain, including language. This was the main point of our article, so on the important points we agree. We will here focus on the points where there is some disagreement between us.

1. How much do we, as a field, know about functional specialization of language-sensitive brain regions?

Although different researchers in the fields of psycho- and neuro-linguistics may vary in how satisfied they are with the progress that has been made in understanding the brain basis of language, we don't think anyone would argue that we have arrived at a complete understanding of what computations each language-sensitive brain region performs and/or what representations it stores/manipulates. So, given that we can probably agree on this, the question becomes: how much do we actually know, as a field, about the brain basis of language in general, and about the questions of functional specialization of language-sensitive brain regions, in particular.

As you point out, we have argued that two fundamental questions concerning the functional architecture of the language system remain unanswered:

• Do language-sensitive brain regions each support one particular aspect of language, or do they support multiple linguistic processes?
• Do language-sensitive regions specialize for (some aspect of) language processing, or do they also support some non-linguistic processes (e.g., general working memory, arithmetic processing, music, action representation, etc.)?

You implied that our pessimistic assessment is due to ignorance. In fact, we are aware of the relevant literature, but we think that the findings that you refer to, among others, either do not pertain to or conclusively answer the questions of functional specificity.

Let us first clarify what we mean by "functional specificity", because perhaps some of the confusion stems from different interpretations of the issues at hand. Here is our working definition: A brain region R is specialized for cognitive function x if this region (i) is engaged in tasks that rely on cognitive function x, and (ii) is not engaged in tasks that do not rely on cognitive function x.

It is important to clarify the difference between a brain region supporting a particular cognitive function and a brain region selectively supporting that cognitive function. A demonstration of a relationship between brain region R and cognitive function x tells us little about R's specificity for

x. For example, various working memory tasks consistently activate dorsolateral prefrontal cortical regions, across studies and labs. However, this does not mean that these regions are selectively engaged in working memory; instead, these regions appear to support a wide range of demanding cognitive tasks (e.g., Duncan, 2001, *Nat Rev Neurosci*; Duncan, 2010, *TiCS*). We elaborate on this issue – with respect to the claim that a portion of Broca's area is selectively engaged in one aspect of syntax (e.g., Grodzinsky & Santi, 2008, *TiCS*) – in our response to Grodzinsky (Fedorenko & Kanwisher, submitted; see also Willems & Hagoort, 2009, *TiCS*, for similar arguments).

Now, let's turn to your claims.

Concerning within-language specificity, you said:
"… there is strong evidence from functional imaging regarding the involvement of distinct cortical regions/circuits in phonemic (STS), lexical-semantic (MTG), prosodic (anterior dorsal STG), and higher-level combinatorial processes (anterior temporal/inferior frontal regions). Additional circuits have been delineated that support auditory-motor integration and auditory/phonological short-term memory. Here are some relevant reviews of this literature: Binder et al., 2000; Hickok and Poeppel, 2007; Indefrey and Levelt, 2004."

Concerning domain specificity of language-sensitive regions, you said:
"Regarding question two, several studies have clearly identified voice-specific responses in the STG (Belin et al. 2000), higher-level speech specific responses in the STS (Scott et al. 2000; Okada et al. 2010), and even what we might as well call the anterior temporal lobe sentence area, given how selective it is to the perception of sentence-level stimulation (Humphries et al., 2006; Humphries et al., 2005; Humphries et al., 2001; Rogalsky and Hickok, 2009; Vandenberghe et al., 2002). But more importantly, some of us have moved beyond the specificity issue with the aim of trying to identify the circuits and computations involved in a given process whether or not it is special to speech."

Then later you pointed out that we have not cited several important findings from the recent literature. Of relevance to the questions of functional specificity are the following (NB: (i) and (iii) are not directly relevant to this question, so we will not discuss those here):
"(ii) convergence on the involvement of the STS in phonemic level processes in speech perception (Leibenthal et al., 2005; Scott & Johnsrude, 2003; Hickok & Poeppel, 2007)";
"(iv) recent progress in mapping the circuit that supports sensory-motor integration in speech processing (Golfinopoulus et al. 2009; Hickok et al., 2009)";
"(v) progress in understanding the basis of hemispheric asymmetries for acoustic and phonemic processing in auditory cortex (Boemio, et al. 2005); Zatorre, et al. 2002)
"(vi) convergence on the idea that anterior temporal regions support some aspect of sentence-level processing (the linguistic equivalent of the FFA)";
"(vii) convergence on the relation between sensory-motor circuits and phonological short-term memory (Buchsbaum et al. 2008; Postle, 2006)".

Although we are familiar with the findings and papers you mention, we don't think that – with respect to any of the brain regions above – we can conclusively say: this region exclusively performs cognitive function x, or cognitive functions x_1-x_n (if a region is truly multifunctional). This is not to say that some of the studies you mention, as well as dozens of other studies, have not

narrowed down the space of possible hypotheses for the function(s) of various regions.

Let's consider one of your examples above: the claim that "anterior temporal regions support some aspect of sentence-level processing"; this area, according to you, is "selective […] to the perception of sentence-level stimulation", and you even go as far as calling it "the linguistic equivalent of the FFA". The papers you cite in support of this claim indeed quite strongly suggest that this region is engaged in processing sentence-level material. However, in order to argue that this region selectively supports sentence-level understanding, more work is needed. In particular, language has been argued to share various properties with a number of non-language tasks (e.g., arithmetic, musical processing, action representation, general executive functions, aspects of social cognition, etc.). Furthermore, tasks in some of these domains have been shown to activate regions in/around anterior temporal cortex (just to provide a few examples: amodal conceptual processing – e.g., Patterson et al., 2007, *Nat Rev Neurosci*; memory for music – e.g., Peretz, 1996, *J of Cog Neurosci*; aspects of social cognition – e.g., Zahn et al., 2007, *PNAS*; Ross & Olson, 2010, *Neuroimage*; face processing – e.g., Sugiura et al., 2001, *Neuroimage*). Perhaps these different non-language tasks activate different portions of the ATL than sentence-level understanding, and perhaps you can dismiss some of these claims / findings on the basis of some design, methodological or analysis flaws. Nevertheless, in order to convince the world that anterior temporal regions are specialized for some aspect of sentence-level processing (as you seem to argue), you would have to show that those regions do not respond to any of the tasks that have been shown to activate regions in/around ATL (as e.g., Kanwisher and colleagues have done for the FFA and other ventral visual stream regions).

Another finding you mention is "the involvement of the STS in phonemic level processes in speech perception". As with the role of the ATL structures in sentence-level understanding, we agree with you that there is strong evidence for the engagement of some portions of STS in phonemic level processes. Perhaps this is even the strongest case of functional specialization in language-related brain regions since across several studies (including the ones you cite) a stronger response to speech sounds has been demonstrated relative to several non-speech auditory control conditions (cf. Leech et al., 2009, *J Neurosci*, for a claim that speech-sound-sensitive areas in STS are not domain specific but rather subserve the processing of acoustically complex sounds more generally and only show higher responses to speech sounds due to more extensive experience with those sounds). However, as with the ATL structures, we think more work is needed. Within language, some studies have argued for the sensitivity of STS regions (including the posterior portion of it) to high-level linguistic manipulations (see e.g., Caplan, 2007, *Language & Linguistics Compass*, for a review of some such studies; see also Friederici et al., 2010, *Hum Brain Mapping*, among others). Furthermore, a number of non-linguistic tasks have been shown to activate regions in/around STS, such as action representation (e.g., Vander Wyk et al., 2009, *Psych Sci*) or aspects of social cognition (e.g., Allison et al., 2000, *TiCS*), among others (see Hein & Knight, 2008, *J Cog Neurosci*, for a recent review). Again, as with the ATL structures, these different tasks may activate regions that are non-overlapping with the speech-sound-sensitive regions. However, this needs to be demonstrated empirically.

In summary, in order to make claims about functional specificity of some region R for a cognitive function x, it is necessary to examine R's response to a wide range of linguistic and non-linguistic tasks that have either been (a) argued to share some properties with x on theoretical grounds, and/or (b) shown to activate cortex in/around R. (Group-based methods can be used for addressing these

questions, but individual-subject analyses are particularly well suited as they will be most likely to uncover dissociations if such are present (Nieto-Castanon, Fedorenko & Kanwisher, in prep.))

Nevertheless, despite the fact that questions of functional specificity have not yet been "nailed" for any of the language-sensitive regions, you are right that previous work has contributed significantly to narrowing down the space of possible hypotheses for the functions of many of these regions (and perhaps this is what you wanted us to acknowledge more in our 2009 paper). Our 2009 paper was not intended to be a complete review of the relevant literature. Instead, our approach in that paper (as well as in the 2010 J Neurophys paper) has been to fairly broadly state the research questions that, in our opinion, remain open as of yet (i.e., the questions of functional specificity of language-sensitive regions). This approach has its limitations: some researchers may feel neglected and/or not properly acknowledged for their contributions to the field. We apologize for any such omissions.

2. Understanding the degree of functional specialization of a region is critical to discerning the computations it supports and/or the representations it stores/manipulates.

You said:
"But more importantly, some of us have moved beyond the specificity issue with the aim of trying to identify the circuits and computations involved in a given process whether or not it is special to speech."

We don't understand this assertion. We also would like to understand the computations that different language-sensitive brain regions perform. However, understanding whether / how specialized a particular region is for language (or a particular aspect of linguistic processing), is essential for understanding the region's computations, so the juxtaposition is confusing.

There are at least three possible research strategies for discerning the nature of the computations/representations in language-sensitive brain regions. One could start by characterizing the language system (i.e., all the language-sensitive brain regions) in terms of their within-language specificity: figuring out whether different regions specialize for particular aspects of language vs. support multiple aspects of language, and understanding how these regions work together to enable language production /comprehension. Then one could use this information to further examine whether these regions support any non-linguistic processes. Alternatively, one could start by characterizing language-sensitive brain regions in terms of their domain specificity: figuring out which, if any, non-linguistic cognitive processes different language-sensitive regions support. Then one could use this information to further examine the precise role of these regions in linguistic processing. Finally, one could simultaneously tackle both questions – (i) within-language specificity, and (ii) domain specificity – and treat the findings with respect to these questions as mutually informative. We are adopting the third approach and hope to use knowledge accumulated in both of these lines of research to draw inferences about the function(s) of each language-sensitive region.

(Of course, fMRI is also not the only method that is needed to answer these questions definitively: high-temporal-resolution methods will be critical for potentially dissociating computations that may be performed by the same cortical tissue but occur at different time-scales.)

Somewhat relatedly, in another one of your posts (http://www.talkingbrains.org/2008/09/brodmann-areas-and-localization-in.html), you said the following:

"Ted Jones provides an instructive reminder in this new age of localization-based neuroscience: "No cortical area is an isolated entity in which a single function is represented. Nor, contrary to many current views, does it merely form one step in a hierarchy of areas proceeding onwards and upwards to some defined or imagined higher function. While there are definite streams of cortico–cortical connections that proceed in identifiable ways from area to area in the cortex, no area is without feedback connections and no area is without re-entrant connections from the thalamus.""

Although it is certainly true that any given brain region is connected to numerous other brain regions, we do not agree with the assertion that "no cortical area is an […] entity in which a single function is represented". Several regions in the ventral visual stream have been shown to be exquisitely specialized for processing visual stimuli of a particular class (see e.g., Kanwisher, 2010, PNAS, for a recent overview). Furthermore, Saxe and colleagues have shown that a region in the right temporo-parietal junction selectively responds to stimuli that require us to think about what another person is thinking (e.g., Saxe & Powell, 2006, Psych Sci, and many other papers; see the publications section on the SaxeLab's website: http://saxelab.mit.edu/publications.php). It is possible that none of the language-sensitive brain regions will show this degree of functional specialization. However, this question deserves to be investigated in the most rigorous way possible given its centrality to understanding human cognition in general.

Given numerous reports of what look like highly specialized linguistic deficits following focal brain damage, we tend to think that at least some degree of functional specialization is bound to be present in the language system. However, we are not on a quest to demonstrate functional specialization of some brain region for language (or for a particular aspect of language). Instead, we want to understand the functional response profiles of different language-sensitive brain regions, and the methods we advocate are perfectly suited for this task. Only by understanding how a region responds to a wide range of cognitive tasks can we begin to make inferences about the computations it performs and the representations that it stores/manipulates.

3. A couple of comments on "methodological schooling".

You said:
"This is frankly a pathetic summary of the state of the field and a pretentious starting point for the methodological schooling that F&K provide in the following sections of their paper."

"Methodological schooling" is not our intention. A few points are worth making here.

a) Many brain imaging methods have a place in investigations of the brain basis of language.

We tried to make this clear in our 2010 paper and in our response to Grodzinsky (Fedorenko & Kanwisher, submitted) and will keep trying to get this point across in our future talks and papers: we strongly believe that individual-subject analyses (both the "basic" subject-specific fROI analyses as well as some more sophisticated individual-subject analyses we've been developing) can go a long way in providing clearer answers to some outstanding questions. However, we are not arguing

for abandoning the traditional, group-based, methods (or for ignoring the knowledge we have gained from those methods). Rather we think that a multi-pronged approach to the study of language – using all the tools available to us – is what's needed for faster progress in the field.

b) How novel are functional localizers in language research?

We do realize and explicitly acknowledge (e.g., Fedorenko et al., 2010, J Neurophys, p. 1178) that some researchers have used individual subject analyses previously (we apologize for not including your studies in this list; we will do that in the future). However, as we note in the paper, to the best of our knowledge, prior to our 2010 paper (and efforts in the clinical literature to localize language-sensitive cortex using fMRI, which have not been adopted in the non-clinical literature), no contrast aimed at localizing language-sensitive brain regions has been validated, i.e., shown to pick out the same regions reliably within and across subjects. Without independent validation of the functional localizer contrasts, it may be difficult to interpret the responses of the ROIs to the critical condition of interest. What are needed are not "ad hoc" localizers that are different across different studies / labs but standardized ones that have been shown to effectively identify the target regions. For example, the localizer that you used in Rogalsky & Hickok (2009, Cereb Cortex) – a contrast between sentences and word lists – identifies only a region in the ATL and does not find any posterior temporal or frontal regions, so does not manage to capture many regions known to be engaged in sentence processing. A key goal of our 2010 J Neurophys paper was to test and establish the robustness of a specific localizer that successfully identifies most of these regions and hence will be of widespread use for many future studies. (Note that although we use the contrast between sentences and nonwords as our main contrast, the contrast between sentences and word lists in our localizer identifies very similar regions to those identified with the sentences > nonword lists contrast, including left frontal and temporal / temporo-parietal regions.)

Another contribution of our work is in developing a method for objectively and efficiently delineating the borders of functional ROIs in individual subjects, the group-constrained subject-specific (GSS) method (for more info see: http://web.mit.edu/evelina9/www/funcloc.html: see Julian et al., in prep., for validation of this method on well characterized high-level visual regions).

c) We invite others to try our methods.

As we mentioned in an email to you, by developing a way to quickly and reliably identify language-sensitive regions, we are not trying to say, "You have all done this wrong; see how we can do it better", but rather, we want to relay something along the lines of, "There may be some limitations in the methods that are currently in use in the vast majority of fMRI studies of language; here is a new approach: let's see how we, as a field, can put our heads together and harness the power of individual subject analyses – by developing a series of standard language localizers and using them consistently across studies and labs – to bring more clarity to where it's needed".

By making all our tools publicly available (http://web.mit.edu/evelina9/www/funcloc.html), we hope that other researchers will join us in adopting the subject-specific approach and that, as a field, we can develop a set of standard localizers for various aspects of language that can be used consistently across studies and labs. We can then – in a joint effort – systematically characterize a set of key brain regions that enable proper linguistic functioning by examining their response to a

wide range of linguistic and non-linguistic stimuli and tasks, thus deriving detailed functional profiles for each region.

We would like to conclude by commenting on the following assertion you make at the end of your post:

"individual subject analyses would be helpful, but it is not a magic bullet (e.g., task selection is more important in my view)"

First, we agree that task selection is highly important in characterizing functional profiles of language-sensitive (or other) regions. However, we don't have to choose between good task design and the use of ROIs: we can do both.

Second, we are advocating individual subject analyses, and we strongly believe that such analyses will bring a clearer understanding of the brain basis of language. We have not shown this yet, so there is no point to argue against this. So far, we developed a method and several new analysis tools that now enable us to pursue questions of functional specificity with rigor that has not so far been possible, because (i) group-based methods are guaranteed to underestimate specificity, and (ii) without a standard localizer used to define ROIs across studies / labs, it is difficult to compare activations across studies / labs, because one can never – with certainty – determine whether an activation peak in one study reflects the activity of the same region as the activation peak in another study, or of a different / nearby region. Of course, it remains to be seen what this method can bring to the field in terms of our understanding of the language architecture. We think it's too early to make pronouncements like the one you make above. Why not give this method a chance? We invite you and others in the field to join our efforts.

Best,
Ev Fedorenko and Nancy Kanwisher

P.S. Our response to Grodzinsky's critique of our 2009 paper will soon be available from Ev's website (http://web.mit.edu/evelina9/www/).

P.P.S. One last comment – not related to the issues above – remains to be made. You made a parallel between the current debate and another debate that I (Ev) was recently involved in: the debate about whether quantitative methods are needed in linguistics research (Gibson & Fedorenko, 2010, *TiCS*; Gibson & Fedorenko, in press, *LCP*). These two debates are quite different, in spite of some superficial similarities. Unlike brain imaging research on language, where a choice of an analysis method may be driven by the nature of the research question and sometimes by researchers' preferences, the use of quantitative methods in linguistics it is not a matter of a research question or preference/opinion. Quantitative methods should be used in linguistics, as they are used in every other branch of cognitive science, and science more generally.

PART 3: August 31, 2010

Greg Hickok wrote:

Hi Ev and Nancy,

First let me thank you for engaging in this discussion. I am thrilled that we are able to debate relevant issues in a publicly accessible format and do so rapidly and in great detail. This kind of exchange of ideas is not afforded by journals (slow and restricted) or conferences (not widely accessible and restricted) and is precisely the reason we launched this blog. It is nice to see it working as intended.

You make a bunch of points so I'll respond piecemeal in separate comments to your post.

The two main points I'd like to make first are: (1) as we discussed in our emails, I think some of the disagreement again stems from a lack of clarity in what you are arguing for in your papers -- things become a bit more clear as you elaborate in less formal venues like this one. (2) I disagree that functional specificity, as it is studied traditionally is a critical question to answer, and in fact I think it is side tracking the field a bit. I think we are going to end up agreeing on both of these points once our definitions are clear, but let's see...

Clarity in your argument

Your paper asks two questions: "(i) Are distinct cortical regions engaged in different aspects of language? (ii) Are regions engaged in language processing specific to the domain of language?" You conclude that, "Neuroimaging has not yet provided clear answers to either question."

Broadly speaking, I take "different aspects of language" to be domains such as phonological, lexical-semantics, syntax, prosody/intonation and processes such as auditory-motor integration, auditory-visual integration, spectro-temporal analysis of speech, stages of speech production, etc. As I pointed out previously and as you seem to agree, there is strong evidence that, *for example*, the STS supports phonological-level processes, more ventral temporal areas support lexical-semantic level processes, whereas the ATL supports some aspect of processing at the sentence level (it responds more to sentences than the same words presented in an unstructured order).

So, are distinct cortical regions engaged in different aspects of language? Yes. You are just plain wrong about that assertion on my interpretation of your question.

But based on your post, you seem to have in mind a different question. For example, you state:

"we don't think anyone would argue that we have arrived at a complete understanding of what computations each language-sensitive brain region performs and/or what representations it stores/manipulates"

Agreed. But this is a different question. This is the hard question: what *computations* are being performed in a region such that it supports one versus another level of language processing. If this is

what you really had in mind for question #1, it wasn't at all clear.

(I'm curious... Is there agreement on what computations the FFA performs?)

Functional specialization

I agree that looking at the response properties of a region for a range of stimulus and task conditions can provide important information about what an area is doing computationally. Have a look at several of my papers on the ATL, Spt, STS, and Broca's area, which have attempted to do just this. But obsessing over specificity can be misleading if you are focused on the wrong level of analysis. For example, there is a tradition of trying to identify the "phoneme specific" area in speech perception. You'll see a bunch of studies that contrast speech with acoustically similar but non-speech stimuli. Regions that responded to both speech and nonspeech were largely ignored as "general acoustic" processing areas. But this misses the computational point: just because a computation isn't specific to speech doesn't mean it isn't a critical stage of processing. Extracting FM sweep duration and direction is important for speech as well as, say, distinguishing animal sounds. It is in this sense that I say some of us have moved beyond the specificity issue. This was probably a case of lack of clarity on my part because I agree that we need to understand specificity of computation, just not specificity of, for lack of a better term, linguistic domain.

Even this can be misleading though if a region participates in different functions (computations). I've used the example of the vocal tract in the past: it is clearly "specialized" for speech and acquired its organization via evolutionary pressure. However, some of the same articulators participate critically in an unrelated function: digestion. This doesn't mean we shouldn't try to characterize the response properties of a region, it just means we shouldn't be dogmatic about how specificity of response relates to specificity of function.

PART 4: August 31, 2010

Ev Fedorenko and Nancy Kanwisher wrote:

Dear Greg,

We thank you for your comments on our response. It looks like we are definitely making progress on agreeing on various key issues, which is great. A few things remain to be clarified.

1. Functional specificity.

You said:
"Broadly speaking, I take "different aspects of language" to be domains such as phonological, lexical-semantics, syntax, prosody/intonation and processes such as auditory-motor integration, auditory-visual integration, spectro-temporal analysis of speech, stages of speech production, etc."

Yes, this is what we have in mind, although the "cuts" among different linguistic operations may end up being along somewhat different lines (as you seem to suggest later in your response), so it's important to keep an open mind in what distinctions we are looking for within linguistic processing.

You said:
"So, are distinct cortical regions engaged in different aspects of language? Yes. You are just plain wrong about that assertion on my interpretation of your question."

It looks like we are still talking about different notions here. There is an important difference between (a) a region being engaged in a particular aspect of language, and (b) a region being *selectively* engaged in a particular aspect of language (cf. V1 being engaged in face processing vs. the OFA/FFA being selectively engaged in face processing). Although there are many claims (backed up by evidence) about relationships between some linguistic functions and some brain regions (e.g., STS - phonological-level processes, ATL - some aspects of sentence-level understanding), there doesn't appear to be consensus on the selective engagement of various regions in different aspects of language. For example, we know that cortex in/around STS is modulated by phonological-level processes, BUT do we know that these same regions are NOT modulated by e.g., lexical-level manipulations or syntactic-level manipulations, or even non-linguistic processes, like social cognition phenomena (which activate nearby, possibly the same, regions)? We don't think such evidence exists, and we think that understanding how selectively any given region supports a particular aspect of linguistic processing is important for understanding what the region does.

Your example of the vocal tract is consistent with us having different notions of "specificity". In particular, you said:
"it [the vocal tract] is clearly "specialized" for speech and acquired its organization via evolutionary pressure. However, some of the same articulators participate critically in an unrelated function: digestion."

By our definition, the vocal tract is NOT specialized for speech. It supports a range of physiological

processes, including non-speech related ones, like breathing and ingestion. So, it is *critically engaged* in speech (much like V1 is critical for face processing), but speech is not its only function.

Of course, this doesn't mean that we are not interested in a brain region if it's not functionally specialized for language or some aspect of language (it seems like this may be one of the remaining points of confusion); so, we would not argue for not studying the vocal tract as a critical organ of speech just because it also supports non-speech functions. In fact, understanding how the vocal tract solves the problems of respiration and ingestion may tell us something important about how it deals with various aspects of speech production. Similarly, with respect to language-sensitive brain regions, observing certain kind of multi-functionality can provide important clues to what the region is doing. That's what I meant when I said that we are not on merely a quest to find specificity, but we think it's critical to understand how specialized each language-sensitive region is, for particular kinds of stimuli / processes, as this will narrow down the space of hypotheses for the region's function(s).

2. How to study a brain region's computations?

You are right in that, ultimately, we would like to understand the computations that each language-sensitive region performs. And indeed this is a hard question. (We are also interested in the representations that are stored / manipulated by each of these regions - an equally hard question.) You said: "... what *computations* are being performed in a region such that it supports one versus another level of language processing. If this is what you really had in mind for question #1, it wasn't at all clear."

We will try to be more explicit about this in the future. However, we are not sure what other goals (besides understanding a region's computations) could drive studies aimed at understanding how functionally specialized a region is.

You also seem to be implying that only question #1 (the question of within-language specificity) is relevant to the question of the nature of the region's computations. In our opinion, understanding both within-language specialization (i.e., specialization for a particular aspect of language vs. supporting multiple linguistic operations) and the extent to which a region is domain specific (i.e., specific to language vs. supporting some non-linguistic processes) are important for making inferences about the region's computations. In fact, within-language specificity and domain specificity are orthogonal, at least theoretically.

In particular, a brain region could be: 1) specialized for a particular aspect of language (e.g., syntactic processing) AND be specific to language (i.e., not support similar - e.g., structure-building - operations in other domains, like music or action planning); 2) specialized for a particular aspect of language AND not be specific to language (i.e., support similar operations in other domains); 3) NOT specialized for a particular aspect of language, instead supporting multiple aspects of linguistic processing (e.g., accessing meanings of words and combining these meanings into more complex meaning representations) AND be specific to language (i.e., not supporting any non-linguistic processes); 4) NOT specialized for a particular aspect of language AND not be specific to language.

This is why we think it is most productive to pursue both questions (within-language specificity and domain specificity) in parallel and treat findings from these two lines of work as mutually informative / constraining. Further, each of the possible outcomes above gives different kinds of clues about the computations that go on in the region.

We'd love to hear your thoughts on other ways - in addition to understanding the functional specialization of a region - to get at a region's computations. What do you have in mind?

Finally, to answer your question about the computations performed by the FFA, sadly there is no clear answer to this question yet. It is clearly involved in both detecting faces and identifying them, and it discriminates between different features of a face, as well as processing faces more holistically. All these facts are useful and constrain an ultimate story, but they do not yet lead to a strong conclusion about the precise computations conducted in the FFA. However, understanding the responses of the FFA to a range of stimuli and tasks has certainly constrained the possible hypotheses about the computations it may support.

Best,

Ev and Nancy

PART 5: September 1, 2010

Greg Hickok wrote:

Part 5.1

Ev and Nancy,

You suggest that the ATL is not up to snuff in terms of sentence-level processing specificity. You note that to determine true specificity, on par with the FFA, we need to assess lots of functions: music, action representation, etc.

There's always more work to do and new findings can change the current theoretical picture in any domain. A quick search on the FFA for example turned up a recent paper in by Hanson & Schmidt titled, "High Resolution Imaging of the Fusiform Face Area (FFA) Using Multivariate Non-linear Classifiers Shows Diagnosticity for Nonface Categories" (Neuroimage. 2010 Aug 21). So one could argue that research on the FFA is equally not up to snuff either because not all possible stimuli, tasks, analyses have been fully assessed.

But this isn't the point. What you were arguing in your paper was that language people should adopt a "new" approach. My point about about the ATL was to illustrate that language scientists have been using this approach (among others) in exactly the way you urge. We *have* tested ATL responses to sentences against a range of other stimulus and task conditions including non-linguistic meaningful sound "events" (Humphries et al. 2001), semantic manipulations, prosodic manipulations, and more recently music (presented at CNS 2008 and currently in review). To date, sentences activate portions of the ATL better than any other stimulus.

So is the ATL is selective for sentence processing? To the extent that it has been tested, yes, but I don't really care. Determining selectivity is not the goal of this research, it is just a clue to what the underlying computation might be, the latter being the real goal of the research.

Part 5.2

Hi Ev and Nancy,

Quite frankly I think you are a bit overly obsessed with the notion of specificity. Understanding the response properties of a region is a good and useful approach to research and determining the degree of specificity is one outcome of this approach. However, embarking on an endless search for specificity, I fear, will only bog us down in "yes it is", "no it isn't" debates with little thought going into the question of what the area is actually doing. For example, in addition to the FFA paper I noted in a previous post, someone just sent me a link to the following paper apropos of your claim for a specialization of right TPJ:

Mitchell, JP (2008) Activity in Right Temporo-Parietal Junction is Not Selective for Theory-of-Mind. Cerebral Cortex. 18:262-271

As you note, after more than a decade of research on the FFA, the poster child of specificity, we still don't know what it is doing, or indeed whether it is even specific to faces!

I think a better approach is to develop a hypothesis regarding the computational function of a region and to test that hypothesis with whatever type of experiment that best assesses the predictions of that hypothesis.

PART 6: September 1, 2010

Ev Fedorenko and Nancy Kanwisher wrote:

Greg,

It appears that we agree that understanding how specialized a brain region is for particular stimuli / processes is informative to the region's computations.

You then say:
"I think a better approach is to develop a hypothesis regarding the computational function of a region and to test that hypothesis with whatever type of experiment that best assesses the predictions of that hypothesis."

We agree that developing and testing hypotheses about a region's function(s) is what's needed, so it looks like we just disagree on the relative importance of understanding the region's specialization for generating and evaluating such hypotheses. We can just leave it at that: not all scientists have to agree on the priority ranking of the "most important questions" and on how to best address them. However, I think we should be able to point out - as we did in our 2009 and 2010 papers - that the specificity questions have not been answered yet with respect to language-sensitive regions (even if you - and maybe others - think that these questions are not the most important questions to address in studying the language system, or cognition in general). In our opinion, specificity is interesting in itself, and it provides a very direct (arguably even necessary) route to the computational questions.

Best,

Ev and Nancy

P.S. Regarding your comment on the FFA, you might want to look at the second section of the supplement to Kanwisher's 2010 PNAS paper (http://web.mit.edu/bcs/nklab/media/pdfs/Kanwisher.PNAS2010.pdf). And regarding your comment on the rTPJ, you might want to read Scholz et al., 2009 (available from Saxelab's website: http://saxelab.mit.edu/publications.php).