The Functional Architecture of Language Comprehension Mechanisms: Fundamental Principles Revealed with fMRI

By:

Idan Blank

M.A., Tel Aviv University (2011)

Submitted to the Department of Brain and Cognitive Sciences
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

August 2016

[September 2016]

Signature redacted

Signature of Author ....
Department of Brain and Cognitive Sciences
August 10, 2016

Signature redacted

Certified by ....
Nancy Kanwisher
Walter A. Rosenblith Professor of Cognitive Neuroscience

Signature redacted

Accepted by ....
Matthew A. Wilson
Sherman Fairchild Professor of Neuroscience and Picower Scholar
Director of Graduate Education for Brain and Cognitive Sciences

[This page intentionally left blank]

Abstract

The functional architecture of language comprehension mechanisms: fundamental principles revealed with fMRI

A key requirement from cognitive models of language comprehension is that they specify the distinct computational mechanisms that are engaged in language processing and the division of linguistic labor across them. Here, I address this requirement from a cognitive neuroscience perspective by employing functional MRI to study the neural implementation of comprehension processes. My experimental approach, unprecedented in studies of language, combines available methods to simultaneously achieve (i) increased functional resolution, via localization of functional brain regions at the single-participant level; (ii) ecological validity, through data-driven, model-free paradigms using naturalistic stimuli; and (iii) statistical rigor, by explicit comparison of functional profiles across regions.

Using this approach, I first contrast two cortical networks engaged in comprehension: one, the "high-level language network", is selectively recruited by linguistic processing but not by other cognitive functions; another, the "multiple-demand network", is recruited across diverse cognitive tasks, both linguistic and non-linguistic. I show that, during naturalistic cognition, each network shows high synchronization amongst its constituent regions, whereas regions across the two networks are functionally dissociated. Thus, these two systems likely play distinct roles in comprehension, which I then characterize by demonstrating that the language network closely tracks linguistic input whereas the multiple-demand network does not. This finding critically constrains the possible contributions of the multiple-demand system to comprehension.

Next, I focus on the high-level language network and examine two current hypotheses about its internal structure. In one study, I find that activity elicited by syntactic processing is not localized to focal language regions but is instead distributed throughout the network, suggesting that syntax is cognitively inseparable from other aspects of language. In another study, I estimate the timescales over which different language regions integrate linguistic information and find that they share a common profile of temporal integration. Therefore, the topographic division of linguistic labor across this network is not organized along distinct integration timescales.

Collectively, these results account for crucial inconsistencies in the literature and challenge common theoretical views. By characterizing the fundamental functional architecture of comprehension mechanisms, these results provide novel insights into the ontology of linguistic mechanisms that give rise to human language.

[This page intentionally left blank]

Table of contents

Part A: language-selective and domain-general mechanisms in language comprehension

Part B: Distributed processing throughout the high-level language network in the left hemisphere

List of figures

List of tables

Introduction

The capacity for language is unique to the human species (Pinker and Jackendoff, 2005) and universal across cultures. It plays a fundamental psychological role: language generates structures that encode meanings (Frege, 1892; De Saussure, 1916/2011) and are efficiently designed to transmit such meanings across individuals (Jaeger and Levy, 2006; Jaeger, 2010; Piantadosi et al., 2011, 2012), consequently influencing their actions (Austin, 1962). Thus, linguistic structures should reflect the thought structures that they transmit, and characteristics of language use should reflect characteristics of human interactions. Language therefore provides a window into both cognitive and social psychology; formulating a theory of language is a critical step towards understanding the human experience.

In this dissertation, I specifically focus on formulating a cognitive theory of language comprehension. A key desideratum for such a theory is that it commits to a functional ontology of "linguistic kinds": it should carve "the language faculty" at its joints by identifying its constituent mechanisms and establish their respective contributions to comprehension. The work presented here addresses this desideratum by asking: what core principles guide the division of linguistic labor across distinct cognitive systems? It is concerned with broad organizational principles that abstract away from any detailed specification of the algorithms and representations employed by different comprehension mechanisms. Namely, I ask (i) whether special-purpose language mechanisms and domain-general cognitive mechanisms support distinct computations in comprehension; (ii) whether sentence structure (i.e., combinatorial syntax) has its own, specialized processor or is instead inseparable from other aspects of language; and (iii) whether linguistic integration over different timescales (e.g., across words, sentences and paragraphs) are implemented by distinct systems or, rather, by a common system. Such principles constitute what I refer to as the "fundamental functional architecture" of language comprehension mechanisms.

To address these questions I take a cognitive neuroscience approach, employing functional MRI (fMRI) to study the neural basis of language and consequently constrain the functional ontology of language comprehension (Blumstein and Amso, 2013; Mather et al., 2013). Specifically, I critically compare the functional profiles of different cortical regions that are recruited during comprehension with the following reasoning: (i) regions that are engaged in comprehension but show a functional dissociation between their respective activity profiles provide evidence for the existence of distinct computational mechanisms; (ii) if a given aspect of comprehension elicits a distributed functional response, then linguistic labor is not ontologically divided based on this aspect; and (iii) if functional responses to different aspects of comprehension recruit overlapping regions, they might be implemented by a common neural circuit and, by extension, share computational resources. Limitations to these inferences are discussed throughout the dissertation.

Neuroimaging studies of language: methodological concerns

The approach outlined above has been implemented by a vast neuroimaging literature over the past two decades and several models detailing the cortical mechanisms of language comprehension have consequently been suggested (e.g., Friederici, 2002; Pulvermüller, 2002; Friederici, 2009; Friederici, 2012a; Price, 2012; Friederici and Gierhan, 2013), each implying an underlying functional ontology of language. Unfortunately, empirical support for these models is weak at best given that the evidence accumulated thus far is plagued with inconsistencies (for an overview, see Fedorenko and Kanwisher, 2009). Therefore, pivotal questions remain unresolved. Whereas some of the controversies are theoretically intriguing (see the Summary chapter), many conflicting findings might be attributed to the wide use of questionable methodologies:

First, studies are prone to drawing spurious inferences at the group level that do not accurately represent the functional architecture of language comprehension in individual brains. Such inferences result from the standard use of traditional analyses, which pool data across participants in a voxel-wise fashion under the assumption that a given voxel corresponds to a single functional unit across individuals. This assumption might hold for primary and secondary sensory areas but is invalid in association cortices where linguistic processing is implemented. In these high-level areas, the precise mapping of functional regions onto micro-anatomical landmarks is highly variable across individuals (Paus et al., 1996; Amunts et al., 1999; Tomaiuolo et al., 1999; Wise et al., 2001; Chein et al., 2002; Juch et al., 2005; Frost and Goebel, 2012; Tahmasebi et al., 2012), such that a voxel can belong to one functional system in one participant but to a different system in another (Poldrack, 2006; Fischl et al., 2008; Fedorenko et al., 2012b). This variability might lead to erroneous inferences by (i) preventing the detection of activity in a functional region due to little spatial overlap across individuals (low sensitivity); or (ii) causing functionally distinct activations that spatially overlap across individuals to appear as originating from a single region (low functional resolution) (Fedorenko and Kanwisher, 2009). Because these inferences do not hold at the individual participant level, they might also partly account for alarming levels of failed replications (see the "reproducibility crisis" in neuroimaging:

Nichols et al., 2016; Poldrack et al., 2016; for a similar concern across psychology, see Open Science Collaboration, 2015).

Second, even when findings are reproducible in close replications of experimental protocols they might still be inconsistent across studies that differ in their operationalization of the same psycholinguistic construct. In other words, convergent validity might be compromised when different tasks or stimuli that are supposed to engage a common linguistic process are instead tapping different processes. For instance, the analysis of sentence structure through syntactic processes has been reported to engage either the inferior frontal and posterior temporal cortices (Stromswold et al., 1996; Cooke et al., 2002; Ben-Shachar et al., 2003; Wartenburger et al., 2003; Constable et al., 2004; Bornkessel et al., 2005; Fiebach et al., 2005; Caplan et al., 2008; Peelle et al., 2009; Meltzer et al., 2010; Christensen et al., 2013) or the anterior temporal lobe (Humphries et al., 2001; Vandenberghe et al., 2002; Humphries et al., 2005; Rogalsky and Hickok, 2009; Baron and Osherson, 2011; Bemis and Pylkkänen, 2011; Zhang and Pylkkanen, 2015), but the two findings rely on different experimental manipulations: the former finding is obtained by contrasting sentences with simpler vs. more complex structures, or grammatical vs. ungrammatical sentences, thus potentially tapping working memory and integration processes (O'Grady, 2011; Gibson et al., 2013) or expectation-based processes (Hale, 2001; Levy, 2008); the latter finding is instead obtained by contrasting stimuli without linguistic structure (e.g., lists of words) to structured sentences, thus potentially tapping integration of word meanings via semantic processes. More generally, differences across studies employing different tasks might also reflect differences in ecological validity, with some measures of language processing "in the lab" being less reflective than others of language processing "in the wild".

A third methodological concern is that findings distinguishing among brain regions based on their respective response characteristics do not replicate because they rely on invalid statistical inferences (Nieuwenhuis et al., 2011). Such regions might fail to exhibit diverging functional profiles upon more rigorous scrutiny. Hence, in addition to low convergent validity at the level of operationalization, statistical inferences might have low divergent validity. For example, studies might find a significant effect in one brain region but a non-significant effect in another region and conclude that these regions have distinct functional roles without directly contrasting their effects to each other (e.g., Lerner et al., 2011). Conversely, studies might find that significant activations are spread throughout a distributed set of regions without explicitly comparing across those regions and, consequently, fail to notice regional differences (e.g., Wilson et al., 2008; Honey et al., 2012).

To conclude, prior neuroimaging studies of language processing have commonly relied on practices that hinder reproducibility and generalizability. Therefore, it is essential that future studies meet three methodological standards: namely, they should (i) take into account the inter-individual variability in the precise anatomical location of functional regions engaged in comprehension; (ii) guarantee that experimental tasks measure what they are intended to measure – i.e., natural language processing; and (iii) establish functional distinctions across brain regions based on explicit tests of such hypotheses. Meeting these standards is a prerequisite to producing a more coherent body of findings that may, in turn, advance us towards a functional ontology of language comprehension mechanisms.

The current studies: methodological innovation

To aid in resolving ongoing debates regarding the functional architecture of linguistic processes, I employ here a synergistic combination of available methods to simultaneously meet the three methodological standards detailed above. First, I define cortical regions of interest that are engaged in comprehension using functional localization in individual participants (Saxe et al., 2006; Fedorenko et al., 2010). In other words, I establish a correspondence of brain regions across individuals based on functional criteria, even when their respective anatomical locations do not show precise spatial overlap. Compared to traditional group-based analyses, this method confers higher sensitivity and functional resolution, which are necessary for establishing an ontology of computationally distinct cortical systems (Nieto-Castañón and Fedorenko, 2012).

Second, in all but one study I use a highly naturalistic paradigm in which participants passively listen to narrated stories. A central advantage of such richly structured, natural stimuli is that they elicit neural responses that are more reliable compared to experimentally controlled input (Hasson et al., 2010). Moreover, natural stimuli enable broad sampling of the space of processes recruited during comprehension. Critically, the latter feature is traded off against the ability to isolate particular linguistic functions, so this method remains largely agnostic as to what aspects of the input drive the neural response. Thus, instead of modeling signal time-courses based on some pre-specified hypotheses, signals are analyzed using model-free, data-driven approaches that compare time-courses across regions, within and between individuals: within individuals, to establish which regions are highly synchronized in their respective time-courses and likely belong to the same functional network or, conversely, which are dissociated in their respective time-courses and likely support distinct computations (Biswal et al., 1995; Yeo et al., 2011); and across individuals, to distinguish between regions that show high inter-individual reliability in their time-courses and thus track the input very closely from regions with less reliable responses that likely track the input less closely (Hasson et al., 2004).

Third, I directly compare functional profiles across regions through statistical tests that license valid inferences regarding inter-regional differences (Nieuwenhuis et al., 2011). To provide evidence that such differences indeed reflect the functional architecture of language comprehension mechanisms, I evaluate whether they go beyond the trivial inter-regional variability in baseline responsiveness (due to, e.g., vascularization: Harrison et al., 2002; Ances et al., 2008; Ekstrom, 2010; Wilson, 2014). Moreover, in cases where different regions engaged in comprehension exhibit indistinguishable functional signatures I demonstrate that the null effect is not due to insufficient power by identifying differences elsewhere in the brain, e.g., between such regions and low-level auditory regions.

Used in conjunction, these methods allow for group-level inferences that accurately reflect individual brains, are ecologically valid and statistically rigorous – an unprecedented scientific standard in neuroimaging studies of language comprehension. Moreover, in all but one study I directly evaluate the reproducibility and/or generalizability of the findings thus obtained by replicating analyses across samples, experimental designs and materials. The work presented here therefore demonstrates the power of this approach to resolve prior inconsistencies in the literature as well as provide novel insights. Importantly, however, this work also suggests some important limits on the kind of discoveries that we could hope to make with neuroimaging.

Towards a functional architecture of language comprehension

Any functional ontology of a particular cognitive domain should look both outward and inward: outward – to understand the place of that domain in the broader architecture of the human mind, and inward – to divide that domain into its constituent components. This dissertation follows these two paths in two separate sections, with Part A looking outward and Part B looking inward. In this section, I briefly describe and motivate each of the questions addressed in each part.

Part A studies the relationship between the language faculty and the rest of high-level cognition by contrasting special-purpose mechanisms that are selectively engaged in linguistic processing with domain-general mechanisms engaged in both linguistic and non-linguistic processes. Historically, identifying aspects of cognition that are implemented by highly specialized mechanisms has been a central approach to articulating a functional ontology of the human mind and has, accordingly, led to major controversies; these were informed by neuropsychological and, later, neuroimaging studies that revealed which cognitive domains receive their own, dedicated neural "real estate" (Finger, 2001; Kanwisher, 2010). In the particular case of language, the question was not only whether specialized mechanisms existed but also, more deeply, whether language shaped thought such that individuals speaking different languages had distinct experiences (De Saussure, 1916/2011; Whorf, 1965/2012; Gleitman and Papafragou, 2005). Here, however, I focus on questions formulated in the opposite direction, concerning the role of general "thought processes" in language comprehension.

The discovery of language-specific mechanisms began with the pioneering discoveries of Paul Broca (1861/2006) and Carl Wernicke (1874/1969) in patients with lesions to perisylvian regions in the left hemisphere. Later studies of people with Aphasia firmly established that language-processing mechanisms could indeed be selectively impaired, with other cognitive mechanisms remaining intact (Luria et al., 1965; Varley and Siegal, 2000; Varley et al., 2005; Apperly et al., 2006; Bek et al., 2010; Willems et al., 2011). Unlike such specialized mechanisms that are necessary for comprehension, many domain-general mechanisms have been suggested to not be strongly required given that they can be impaired without affecting comprehension, e.g., in individuals with savant or Turner syndromes (Netley and Rovet, 1982; Smith et al., 2010). Nonetheless, the latter still contribute to language processing in neurotypical individuals, as indicated by behavioral studies (Fedorenko et al., 2006; 2007; 2009).

In line with these results, some fMRI studies have reported that a system of frontal and temporal regions, primarily in the left hemisphere, is selectively engaged in high-level linguistic processing (Binder et al., 1997; Fedorenko et al., 2011) but not in other, high-level cognitive processes (Monti et al., 2009; Fedorenko et al., 2012b; Monti and Osherson, 2012; Monti et al., 2012). In contrast, a different, bilateral system of frontal, temporal, cingular and opercular regions is recruited in a domain-general manner by both linguistic processing and numerous other cognitive processes (Duncan and Owen, 2000; Duncan, 2010; Fedorenko et al., 2013). Nevertheless, other studies have not supported the distinction between language-specific and domain-general systems, suggesting instead that they functionally overlap (Dehaene et al., 1999; Stanescu-Cosson et al., 2000; Maess et al., 2001; Kaan and Swaab, 2002; Koelsch et al., 2002; Thompson-Schill et al., 2005; Koechlin and Jubault, 2006; Hein and Knight, 2008; Blumstein, 2009; January et al., 2009; Blumstein and Amso, 2013).

Even if we grant that distinct cortical systems implement language-specific vs. domain-general processes, it remains unclear how these systems work together to support on-line language processing under natural conditions, where a wide array of comprehension processes take place. On the one hand, these systems could frequently interact, transferring information back and forth across their respective regions. On the other hand, they could instead be functionally dissociated, each supporting cognitively separable computations that unfold quite independently of one another. To address this question, Chapter 1 (published in Blank et al., 2014) measures the extent of functional integration across the two systems both during story comprehension – when their activity is evoked by linguistic input, and during mind wandering ("rest") – when such activity is not externally driven.

Chapter 2 then characterizes the respective contributions of these two systems to comprehension. Specifically, because activity in each system is sensitive to some linguistic features, it has been implicitly assumed that both systems closely track linguistic input (see also Wilson et al., 2008; Honey et al., 2012). However, it is possible that only one of these systems exhibits such close tracking, with the other being recruited for linguistic processing less frequently and only in particular circumstances. Therefore, I critically assess the assumption of close input tracking by measuring the reliability of neural activity in each of the two systems during story comprehension.

Part B then focuses on the language-specific cortical system and examines the division of linguistic labor across its constituent regions, an active area of investigation and disagreement in the neuroimaging literature. Specifically, whereas some studies claim that specific language regions are specialized for particular processes (for reviews, see Friederici, 2002; Price, 2012), others report distributed and overlapping activations across different processes (Fedorenko et al., 2010; Bedny et al., 2011; Fedorenko et al., 2012a; Bautista and Wilson, 2016). Therefore, here I evaluate two central hypotheses regarding the internal organization of the language system:

In Chapter 3 (published in Blank et al., 2016), I ask whether the processing of sentence structure – i.e., syntax – is localized to particular regions within the language system and thus ontologically distinct, perhaps, from other linguistic processes. Syntactic processing has a prominent role in psycholinguistic theories, because it is the most important formal tool for combining a finite inventory of linguistic units (i.e., words) into infinitely many possible expressions. Moreover, studies of syntactic processing by Chomsky (1956), in large part, directly led to the cognitive revolution in psychology (Gardner, 2008). Those studies separated syntactic processing from the other critical aspect of comprehension – the mental lexicon, claiming that parsing sentences according to knowledge about how words are allowed to combine is fundamentally different from retrieving word meanings. This separation fit with the intuition that we can parse grammatical sentences even if they are nonsensical, such as Louis Carroll's Jabberwocky poem (1872/2001).

Following early theories, many neuroimaging studies searched for the "seat of syntax" in the brain and suggested that Broca's region was specialized for syntactic processing (Ullman, 2004; Bornkessel and Schlesewsky, 2006; Grodzinsky and Friederici, 2006; Grodzinsky and

Santi, 2008; Friederici, 2009; Baggio and Hagoort, 2011; Friederici, 2011; Tyler et al., 2011; Friederici, 2012b; Duffau et al., 2014). Many studies have also reported focal activations for syntactic processing in posterior temporal regions (for a meta-analysis see Friederici, 2011). However, neuropsychological studies found that syntactic deficits resulted from damage to any of a distributed set of regions, suggesting instead that syntactic processing engages the entire language system (Caplan et al., 1996; Dick et al., 2001; Wilson and Saygin, 2004; Mesulam et al., 2015). These latter findings are consistent with current proposals that no longer advocate for the traditional separation between "words" and "rules" (Joshi et al., 1975; Bresnan, 1982; Schabes et al., 1988; Pollard and Sag, 1994; Chomsky, 1995; Goldberg, 1995; Bybee, 1998, 2010; Jackendoff, 2002, 2007; Culicover and Jackendoff, 2005).

To reconcile these conflicting findings, I propose that distributed syntactic processing could erroneously appear localized to particular regions under certain analysis regimes. I then evaluate this account by first measuring the sensitivity of different language regions to manipulations of syntactic complexity in a commonly used task, and then applying analysis methods that differ in sensitivity and rigor. This experiment does not employ a natural language comprehension paradigm, but it therefore clearly motivates the use of such a paradigm in Chapter 4, where I attempt to find a different division of labor across language regions.

In this last chapter, I test whether language comprehension across high-level language regions is topographically organized according to the timescale of information integration, creating a cortical hierarchy along which the linguistic input is integrated over increasingly longer time-windows. The grain size of temporal integration is an intuitively appealing candidate dimension for organizing the comprehension of spoken language: acoustic input is integrated into phonemes, which in turn form morphemes and words that can combine into phrases and sentences, the latter then giving rise to paragraphs and finally full narratives. Indeed, early linguistic theories treated representations of different grain sizes as the "natural kinds" of language. However, contemporary accounts postulate a flexible representational continuum crossing the traditional boundaries between sounds, words and sentences (see Goldberg, 1995, for a leading psycholinguistic framework), consistent with behavioral evidence that comprehenders keep track of information that straddles such boundaries (Clifton et al., 1984; MacDonald et al., 1994; Trueswell et al., 1994; Garnsey et al., 1997; Traxler et al., 2002; Farmer et al., 2006; Reali and Christiansen, 2007; Gennari and MacDonald, 2008; Schmidtke et al., 2014).

Surprisingly, recent neuroimaging evidence might support earlier theories, with a cortical hierarchy of temporal integration windows having recently been reported (Lerner et al., 2011). Still, even though the topography of this hierarchy seems to spatially overlap with the high-level language system, direct functional correspondence between the two has not been established. Therefore, here I suggest that different language regions in fact all share a common functional profile within this hierarchy, in line with current psycholinguistic proposals. As in the previous chapter, I propose and test an account of how analyses with insufficient sensitivity and rigor might result in inaccurate characterization of language regions, such that they appear to exhibit distinct cortical integration timescales.

The cognitive horizon for neuroimaging studies of language

The questions addressed by this work together with the neuroimaging tools used to investigate them are a promising scientific compass towards a functional characterization of human language across multiple levels of analysis (Marr, 1982). The implications for an implementation level understanding of language processing by neural circuits are self-evident. The implications for an algorithm-and-representation level understanding of language processing are also clear; determining the fundamental functional architecture of human language by dividing linguistic labor into ontological "kinds" is a precondition for specifying the form and content of their respective inputs and outputs or the associated input-output transformations. Some of the work reported here already importantly constrains the space of possible algorithms and representations employed by different systems that are engaged in comprehension.

Finally, such a research enterprise also has implications for a computational level understanding of language processing, i.e., explaining why language is designed the way it is and what goals this design is optimized for achieving. To illustrate this claim, consider first that whatever such goals might be, they appear to not be entirely met by other high-level cognitive functions, given that language processing partly depends on special-purpose mechanisms (Newport, 1990; Fedorenko and Varley, 2016). Indeed, psycholinguistic theories attempt to characterize the design of human language by invoking organizational principles that are unique to linguistic processing (e.g., Chomsky, 1995; Goldberg, 1995; Jackendoff, 2002; Culicover and Jackendoff, 2005; Bybee, 2010). Nonetheless, some aspects of this design might simply result from domain-general constraints: language evolved to fit the pre-existing neural hardware (Christiansen and Chater, 2008) and general cognitive architecture, so some of its functional characteristics may reflect constraints imposed by the structure of, e.g., working memory processes (Gibson, 1998; Van Dyke and Johns, 2012) or processes of prediction (Hale, 2001; Levy, 2008, 2013). More generally, an emerging tradition in computational linguistics suggests that language processing can be modeled using the same tools that account for much of high-level cognition (namely, probabilistic inferences over symbolic representations: Chater and Manning, 2006; Goodman and Lassiter, 2014). To date, special-purpose and domain-general computational notions proceed to evolve quite independently, and integrating them remains challenging. How is the design of human language trades-off language-specific goals and domain-general constraints? Which principles of this design are unique to language processing, which have homologues elsewhere in the human mind, and which simply reflect the workings of general mechanisms shared across domains (e.g., Hauser et al., 2002; Jackendoff and Pinker, 2005; Pinker and Jackendoff, 2005)?

A computational level understanding of human language also concerns the way language processing deals with multiple sources of linguistic information. Linguistic communication conveys meaning through multiple forms, including (but not limited to) morphemes, words, phrases and sentences, prosody, and non-literal, pragmatic devices. If some of these forms are designed to achieve distinct computational goals, then what are these goals? If human language is designed to process some of these forms along distinct processing streams, then what is the advantage of such design?

These questions could perhaps be more accurately answered once a functional ontology of linguistic mechanisms is established: if we are to ask what the design of human language is optimized for, we could benefit from knowing what the design of human language is. The studies that follow demonstrate that by neuroimaging the functional architecture of language comprehension we can take initial steps towards outlining an answer to such questions. First, this work contributes to our understanding of the relationship between language-specific and domain-general resources in linguistic processing – i.e., how frequently they interact and what their respective functional roles are (Fedorenko, 2014). Second, this work evaluates suggestions for the dimensions along which linguistic processing might be organized. Altogether, such work might allow us to "reverse engineer" what goals human language is designed to achieve and what it was evolved for (Christiansen and Kirby, 2003).

These theoretical implications are cognitive horizons for the neuroimaging of human language processing. They might be far, yet based on the studies presented here I believe that they are not out of reach. Let us journey on.

A note on the structure of this dissertation

This dissertation is designed as a collection of journal papers. Thus, chapters differ in the organization of different sections (e.g., Chapter 2 does not have a separate Introduction section, and the Materials and Methods section follows the Discussion). Some chapters provide an exhaustive (perhaps, too exhaustive) investigation of the results from multiple perspectives, usually to address concerns raised by reviewers; others are relatively succinct, in line with the constraints on word count in some journals. Because each chapter was independently written, a noticeable degree of repetition exists across them in methodological description, the overview of prior literature and some of the theoretical points discussed. I apologize for any inconvenience these might cause in reading the dissertation as a unified piece.

References

Amunts K, Schleicher A, Bürgel U, Mohlberg H, Uylings H, Zilles K (1999) Broca's region revisited: Cytoarchitecture and intersubject variability. Journal of Comparative Neurology 412:319-341.

Ances BM, Leontiev O, Perthen JE, Liang C, Lansing AE, Buxton RB (2008) Regional differences in the coupling of cerebral blood flow and oxygen metabolism changes in response to activation: Implications for BOLD-fMRI. NeuroImage 39:1510-1521.

Apperly IA, Samson D, Carroll N, Hussain S, Humphreys G (2006) Intact first-and second-order false belief reasoning in a patient with severely impaired grammar. Social Neuroscience 1:334-348.

Austin JL (1962) How to Do Things with Words. Oxford, UK: Oxford University Press.

Baggio G, Hagoort P (2011) The balance between memory and unification in semantics: A dynamic account of the N400. Language and Cognitive Processes 26:1338-1367.

Baron SG, Osherson D (2011) Evidence for conceptual combination in the left anterior temporal lobe. NeuroImage 55:1847-1852.

Bautista A, Wilson SM (2016) Neural responses to grammatically and lexically degraded speech. Language, Cognition and Neuroscience 31:567-574.

Bedny M, Pascual-Leone A, Dodell-Feder D, Fedorenko E, Saxe R (2011) Language processing in the occipital cortex of congenitally blind adults. Proceedings of the National Academy of Sciences 108:4429-4434.

Bek J, Blades M, Siegal M, Varley RA (2010) Language and spatial reorientation: Evidence from severe aphasia. Journal of Experimental Psychology: Learning, Memory, and Cognition 36:646-658.

Bemis DK, Pylkkänen L (2011) Simple composition: A magnetoencephalography investigation into the comprehension of minimal linguistic phrases. The Journal of Neuroscience 31:2801-2814.

Ben-Shachar M, Hendler T, Kahn I, Ben-Bashat D, Grodzinsky Y (2003) The neural reality of syntactic transformations evidence from functional magnetic resonance imaging. Psychological Science 14:433-440.

Binder JR, Frost JA, Hammeke TA, Cox RW, Rao SM, Prieto T (1997) Human brain language areas identified by functional magnetic resonance imaging. The Journal of Neuroscience 17:353-362.

Biswal B, Zerrin Yetkin F, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo - planar MRI. Magnetic Resonance in Medicine 34:537-541.

Blank I, Kanwisher N, Fedorenko E (2014) A functional dissociation between language and multiple-demand systems revealed in patterns of BOLD signal fluctuations. Journal of Neurophysiology 112:1105-1118.

Blank I, Balewski Z, Mahowald K, Fedorenko E (2016) Syntactic processing is distributed across the language system. NeuroImage 127:307-323.

Blumstein SE (2009) Reflections on the cognitive neuroscience of language. In: The Cognitive Neurosciences (Gazzaniga M, ed), pp 1235-1240. Cambridge, MA: MIT Press.

Blumstein SE, Amso D (2013) Dynamic functional organization of language: Insights from functional neuroimaging. Perspect Psychol Sci 8:44-48.

Bornkessel I, Schlesewsky M (2006) The extended argument dependency model: A neurocognitive approach to sentence comprehension across languages. Psychological Review 113:787.

Bornkessel I, Zysset S, Friederici AD, von Cramon DY, Schlesewsky M (2005) Who did what to whom? The neural basis of argument hierarchies during language comprehension. NeuroImage 26:221-233.

Bresnan J (1982) The Mental Representation of Grammatical Relations. Cambridge, MA: MIT Press.

Broca P (1861/2006) Comments regarding the seat of the faculty of spoken language, followed by an observation of aphemia (loss of speech). In: Broca's Region (Grodzinsky Y, Amunts K, eds), pp 291-304. New York, NY: Oxford University Press.

Bybee J (1998) A functionalist approach to grammar and its evolution. Evolution of Communication 2:249-278.

Bybee J (2010) Language, Usage and Cognition. Cambridge, UK: Cambridge University Press.

Caplan D, Hildebrandt N, Makris N (1996) Location of lesions in stroke patients with deficits in syntactic processing in sentence comprehension. Brain 119:933-949.

Caplan D, Stanczak L, Waters G (2008) Syntactic and thematic constraint effects on blood oxygenation level dependent signal correlates of comprehension of relative clauses. Journal of Cognitive Neuroscience 20:643-656.

Carroll L (1872/2001) Jabberwocky. In: Jabberwocky and Other Poems, p 17. Mineola, NY: Dover Publications.

Chater N, Manning CD (2006) Probabilistic models of language processing and acquisition. Trends in Cognitive Sciences 10:335-344.

Chein JM, Fissell K, Jacobs S, Fiez JA (2002) Functional heterogeneity within Broca's area during verbal working memory. Physiology and Behavior 77:635-639.

Chomsky N (1956) Three models for the description of language. IRE Transactions on Information Theory 2:113-124.

Chomsky N (1995) The Minimalist Program. Cambridge, MA: MIT Press.

Christensen KR, Kizach J, Nyvad AM (2013) The processing of syntactic islands – an fMRI study. Journal of Neurolinguistics 26:239-251.

Christiansen MH, Kirby S (2003) Language Evolution. Oxford, UK: Oxford University Press.

Christiansen MH, Chater N (2008) Language as shaped by the brain. Behavioral and Brain Sciences 31:489-509.

Clifton C, Frazier L, Connine C (1984) Lexical expectations in sentence comprehension. Journal of Verbal Learning and Verbal Behavior 23:696-708.

Collaboration OS (2015) Estimating the reproducibility of psychological science. Science 349:aac4716.

Constable RT, Pugh KR, Berroya E, Mencl WE, Westerveld M, Ni W, Shankweiler D (2004) Sentence complexity and input modality effects in sentence comprehension: an fMRI study. NeuroImage 22:11-21.

Cooke A, Zurif EB, DeVita C, Alsop D, Koenig P, Detre J, Gee J, Pinango M, Balogh J, Grossman M (2002) Neural basis for sentence comprehension: Grammatical and short - term memory components. Human Brain Mapping 15:80-94.

Culicover PW, Jackendoff R (2005) Simpler Syntax. Oxford, UK: Oxford University Press.

De Saussure F (1916/2011) Course in General Linguistics. New York, NY: Columbia University Press.

Dehaene S, Spelke E, Pinel P, Stanescu R, Tsivkin S (1999) Sources of mathematical thinking: Behavioral and brain-imaging evidence. Science 284:970-974.

Dick F, Bates E, Wulfeck B, Utman JA, Dronkers N, Gernsbacher MA (2001) Language deficits, localization, and grammar: Evidence for a distributive model of language breakdown in aphasic patients and neurologically intact individuals. Psychological Review 108:759-788.

Duffau H, Moritz-Gasser S, Mandonnet E (2014) A re-examination of neural basis of language processing: Proposal of a dynamic hodotopical model from data provided by brain stimulation mapping during picture naming. Brain and Language 131:1-10.

Duncan J (2010) The multiple-demand (MD) system of the primate brain: Mental programs for intelligent behaviour. Trends in Cognitive Sciences 14:172-179.

Duncan J, Owen AM (2000) Common regions of the human frontal lobe recruited by diverse cognitive demands. Trends in Neurosciences 23:475-483.

Ekstrom A (2010) How and when the fMRI BOLD signal relates to underlying neural activity: The danger in dissociation. Brain Research Reviews 62:233-244.

Farmer TA, Christiansen MH, Monaghan P (2006) Phonological typicality influences on-line sentence comprehension. Proceedings of the National Academy of Sciences 103:12203-12208.

Fedorenko E (2014) The role of domain-general cognitive control in language comprehension. Frontiers in Psychology 5:335.

Fedorenko E, Kanwisher N (2009) Neuroimaging of language: Why hasn't a clearer picture emerged? Language and Linguistics Compass 3:839-865.

Fedorenko E, Varley R (2016) Language and thought are not the same thing: Evidence from neuroimaging and neurological patients. Annals of the New York Academy of Sciences 1369:132-153.

Fedorenko E, Gibson E, Rohde D (2006) The nature of working memory capacity in sentence comprehension: Evidence against domain-specific working memory resources. Journal of Memory and Language 54:541-553.

Fedorenko E, Gibson E, Rohde D (2007) The nature of working memory in linguistic, arithmetic and spatial integration processes. Journal of Memory and Language 56:246-269.

Fedorenko E, Behr MK, Kanwisher N (2011) Functional specificity for high-level linguistic processing in the human brain. Proceedings of the National Academy of Sciences 108:16428-16433.

Fedorenko E, Nieto-Castañón A, Kanwisher N (2012a) Lexical and syntactic representations in the brain: an fMRI investigation with multi-voxel pattern analyses. Neuropsychologia 50:499-513.

Fedorenko E, Duncan J, Kanwisher N (2012b) Language-selective and domain-general regions lie side by side within Broca's area. Current Biology 22:2059-2062.

Fedorenko E, Duncan J, Kanwisher N (2013) Broad domain generality in focal regions of frontal and parietal cortex. Proceedings of the National Academy of Sciences 110:16616-16621.

Fedorenko E, Patel A, Casasanto D, Winawer J, Gibson E (2009) Structural integration in language and music: Evidence for a shared system. Memory and Cognition 37:1-9.

Fedorenko E, Hsieh P-J, Nieto-Castañón A, Whitfield-Gabrieli S, Kanwisher N (2010) New method for fMRI investigations of language: Defining ROIs functionally in individual subjects. Journal of Neurophysiology 104:1177-1194.

Fiebach CJ, Schlesewsky M, Lohmann G, Von Cramon DY, Friederici AD (2005) Revisiting the role of Broca's area in sentence processing: Syntactic integration versus syntactic working memory. Human Brain Mapping 24:79-91.

Finger S (2001) Origins of Neuroscience: a History of Explorations into Brain Function. New York, NY: Oxford University Press.

Fischl B, Rajendran N, Busa E, Augustinack J, Hinds O, Yeo BTT, Mohlberg H, Amunts K, Zilles K (2008) Cortical folding patterns and predicting cytoarchitecture. Cerebral Cortex 18:1973-1980.

Frege G (1892) On sense and reference. In: Translations from the Philosophical Writings of Gottlob Frege (Peter G, Max B, eds). Oxford, UK: Blackwell Publishers.

Friederici AD (2002) Towards a neural basis of auditory sentence processing. Trends in Cognitive Sciences 6:78-84.

Friederici AD (2009) Pathways to language: Fiber tracts in the human brain. Trends in Cognitive Sciences 13:175-181.

Friederici AD (2011) The brain basis of language processing: From structure to function. Physiological Reviews 91:1357-1392.

Friederici AD (2012a) Language comprehension: A biological perspective: Springer Science & Business Media.

Friederici AD (2012b) The cortical language circuit: From auditory perception to sentence comprehension. Trends in Cognitive Sciences 16:262-268.

Friederici AD, Gierhan SM (2013) The language network. Current Opinion in Neurobiology 23:250-254.

Frost MA, Goebel R (2012) Measuring structural–functional correspondence: Spatial variability of specialised brain regions after macro-anatomical alignment. NeuroImage 59:1369-1381.

Gardner H (2008) The Mind's New Science: A History of the Cognitive Revolution. New York, NY: Basic Books.

Garnsey SM, Pearlmutter NJ, Myers E, Lotocky MA (1997) The contributions of verb bias and plausibility to the comprehension of temporarily ambiguous sentences. Journal of Memory and Language 37:58-93.

Gennari SP, MacDonald MC (2008) Semantic indeterminacy in object relative clauses. Journal of Memory and Language 58:161-187.

Gibson E (1998) Linguistic complexity: Locality of syntactic dependencies. Cognition 68:1-76.

Gibson E, Tily H, Fedorenko E (2013) The processing complexity of English relative clauses. In: Language Down the Garden Path: The Cognitive and Biological Basis for Linguistic Structure (Sanz M, Laka I, Tanenhaus MK, eds). Oxford, UK: Oxford University Press.

Gleitman L, Papafragou A (2005) Language and thought. In: Cambridge Handbook of Thinking and Reasoning (Holyoak KJ, Morrison RG, eds), pp 633-661. New York, NY: Cambridge University Press.

Goldberg AE (1995) Constructions: A construction grammar approach to argument structure. Chicago, IL: University of Chicago Press.

Goodman ND, Lassiter D (2014) Probabilistic semantics and pragmatics: Uncertainty in language and thought. In: Handbook of Contemporary Semantic Theory (Lappin S, Fox C, eds). Malden, MA: Wiley-Blackwell.

Grodzinsky Y, Friederici AD (2006) Neuroimaging of syntax and syntactic processing. Current Opinion in Neurobiology 16:240-246.

Grodzinsky Y, Santi A (2008) The battle for Broca's region. Trends in Cognitive Sciences 12:474-480.

Hale J (2001) A probabilistic earley parser as a psycholinguistic model. In: Proceedings of the second meeting of the North American Chapter of the Association for Computational Linguistics on Language technologies, pp 1-8: Association for Computational Linguistics.

Harrison RV, Harel N, Panesar J, Mount RJ (2002) Blood capillary distribution correlates with hemodynamic-based functional imaging in cerebral cortex. Cerebral Cortex 12:225-233.

Hasson U, Malach R, Heeger DJ (2010) Reliability of cortical activity during natural stimulation. Trends in Cognitive Sciences 14:40-48.

Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R (2004) Intersubject synchronization of cortical activity during natural vision. Science 303:1634-1640.

Hauser MD, Chomsky N, Fitch WT (2002) The faculty of language: What is it, who has it, and how did it evolve? Science 298:1569-1579.

Hein G, Knight RT (2008) Superior temporal sulcus—it's my area; or is it? Journal of Cognitive Neuroscience 20:2125-2136.

Honey CJ, Thompson CR, Lerner Y, Hasson U (2012) Not lost in translation: Neural responses shared across languages. The Journal of Neuroscience 32:15277-15283.

Humphries C, Willard K, Buchsbaum B, Hickok G (2001) Role of anterior temporal cortex in auditory sentence comprehension: an fMRI study. Neuroreport 12:1749-1752.

Humphries C, Love T, Swinney D, Hickok G (2005) Response of anterior temporal cortex to syntactic and prosodic manipulations during sentence processing. Human Brain Mapping 26:128-138.

Jackendoff R (2002) Foundation of Language: Brain, Meaning, Grammar, Evolution. Oxford, UK: Oxford University Press.

Jackendoff R (2007) A parallel architecture perspective on language processing. Brain Research 1146:2-22.

Jackendoff R, Pinker S (2005) The nature of the language faculty and its implications for evolution of language (Reply to Fitch, Hauser, and Chomsky). Cognition 97:211-225.

Jaeger TF (2010) Redundancy and reduction: Speakers manage syntactic information density. Cognitive Psychol 61:23-62.

Jaeger TF, Levy RP (2006) Speakers optimize information density through syntactic reduction. In: Advances in Neural Information Processing Systems (Schlökopf B, Platt J, Hoffman T, eds), pp 849-856. Cambridge, MA: MIT Press.

January D, Trueswell JC, Thompson-Schill SL (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 21:2434-2444.

Joshi AK, Levy LS, Takahashi M (1975) Tree adjunct grammars. Journal of Computer and System Sciences 10:136-163.

Juch H, Zimine I, Seghier ML, Lazeyras F, Fasel JH (2005) Anatomical variability of the lateral frontal lobe surface: Implication for intersubject variability in language neuroimaging. NeuroImage 24:504-514.

Kaan E, Swaab TY (2002) The brain circuitry of syntactic comprehension. Trends in Cognitive Sciences 6:350-356.

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 107:11163-11170.

Koechlin E, Jubault T (2006) Broca's area and the hierarchical organization of human behavior. Neuron 50:963-974.

Koelsch S, Gunter TC, Zysset S, Lohmann G, Friederici AD (2002) Bach speaks: A cortical "language-network" serves the processing of music. NeuroImage 17:956-966.

Lerner Y, Honey CJ, Silbert LJ, Hasson U (2011) Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. The Journal of Neuroscience 31:2906-2915.

Levy RP (2008) Expectation-based syntactic comprehension. Cognition 106:1126-1177.

Levy RP (2013) Memory and surprisal in human sentence comprehension. Sentence Processing:78-114.

Luria AR, Tsvetkova LS, Futer DS (1965) Aphasia in a composer. Journal of the Neurological Sciences 2:288-292.

MacDonald MC, Pearlmutter NJ, Seidenberg MS (1994) The lexical nature of syntactic ambiguity resolution. Psychological Review 101:676-703.

Maess B, Koelsch S, Gunter TC, Friederici AD (2001) Musical syntax is processed in Broca's area: an MEG study. Nature Neuroscience 4:540-545.

Marr D (1982) The philosophy and the approach. In: Vision: A Computational Investigation into the Human Representation and Processing of Visual Information, pp 8-38. Cambridge, MA: MIT Press.

Mather M, Cacioppo JT, Kanwisher N (2013) How fMRI can inform cognitive theories. Perspect Psychol Sci 8:108-113.

Meltzer JA, McArdle JJ, Schafer RJ, Braun AR (2010) Neural aspects of sentence comprehension: Syntactic complexity, reversibility, and reanalysis. Cerebral Cortex 20:1853-1864.

Mesulam MM, Thompson CK, Weintraub S, Rogalski EJ (2015) The Wernicke conundrum and the anatomy of language comprehension in primary progressive aphasia. Brain 138:2423-2437.

Monti MM, Osherson DN (2012) Logic, language and the brain. Brain Research 1428:33-42.

Monti MM, Parsons LM, Osherson DN (2009) The boundaries of language and thought in deductive inference. Proceedings of the National Academy of Sciences 106:12554-12559.

Monti MM, Parsons LM, Osherson DN (2012) Thought beyond language: neural dissociation of algebra and natural language. Psychological Science 23:914-922.

Netley CT, Rovet J (1982) Atypical hemispheric lateralization in Turner syndrome subjects. Cortex 18:377-384.

Newport EL (1990) Maturational constraints on language learning. Cognitive Science 14:11-28.

Nichols TE, Das S, Eickhoff SB, Evans AC, Glatard T, Hanke M, Kriegeskorte N, Milham MP, Poldrack RA, Poline J-B (2016) Best practices in data analysis and sharing in meuroimaging using MRI. bioRxiv:054262.

Nieto-Castañón A, Fedorenko E (2012) Subject-specific functional localizers increase sensitivity and functional resolution of multi-subject analyses. NeuroImage 63:1646-1669.

Nieuwenhuis S, Forstmann BU, Wagenmakers E-J (2011) Erroneous analyses of interactions in neuroscience: A problem of significance. Nature Neuroscience 14:1105-1107.

O'Grady W (2011) Relative clauses: Processing and acquisition. In: The Acquisition of Relative Clauses: Processing, Typology and Function (Kidd E, ed), pp 13-38. Philadelphia, PA: John Benjamin Publishing.

Paus T, Tomaiuolo F, Otaky N, MacDonald D, Petrides M, Atlas J, Morris R, Evans AC (1996) Human cingulate and paracingulate sulci: Pattern, variability, asymmetry, and probabilistic map. Cerebral Cortex 6:207-214.

Peelle JE, Troiani V, Wingfield A, Grossman M (2009) Neural processing during older adults' comprehension of spoken sentences: Age differences in resource allocation and connectivity. Cerebral Cortex 20:bhp142.

Piantadosi ST, Tily H, Gibson E (2011) Word lengths are optimized for efficient communication. Proceedings of the National Academy of Sciences 108:3526-3529.

Piantadosi ST, Tily H, Gibson E (2012) The communicative function of ambiguity in language. Cognition 122:280-291.

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

Poldrack R, Baker CI, Durnez J, Gorgolewski K, Matthews PM, Munafo M, Nichols T, Poline J-B, Vul E, Yarkoni T (2016) Scanning the horizon: Future challenges for neuroimaging research. bioRxiv:059188.

Poldrack RA (2006) Can cognitive processes be inferred from neuroimaging data? Trends in Cognitive Sciences 10:59-63.

Pollard C, Sag IA (1994) Head-driven phrase structure grammar. Chicago, IL: University of Chicago Press.

Price CJ (2012) A review and synthesis of the first 20 years of PET and fMRI studies of heard speech, spoken language and reading. NeuroImage 62:816-847.

Pulvermüller F (2002) The neuroscience of language: On brain circuits of words and serial order: Cambridge University Press.

Reali F, Christiansen MH (2007) Processing of relative clauses is made easier by frequency of occurrence. Journal of Memory and Language 57:1-23.

Rogalsky C, Hickok G (2009) Selective attention to semantic and syntactic features modulates sentence processing networks in anterior temporal cortex. Cerebral Cortex 19:786-796.

Saxe R, Brett M, Kanwisher N (2006) Divide and conquer: A defense of functional localizers. NeuroImage 30:1088-1096.

Schabes Y, Abeille A, Joshi AK (1988) Parsing strategies with "lexicalized" grammars: Application to tree adjoining grammars. In: Proceedings of the 12th Conference on Computational Linguistics, pp 578-583. Stroudsburg, PA.

Schmidtke D, Conrad M, Jacobs AM (2014) Phonological iconicity. Frontiers in Psychology 5:80.

Smith N, Tsimpli I, Morgan G, Woll B (2010) The signs of a savant: Language against the odds. Cambridge, UK: Cambridge University Press.

Stanescu-Cosson R, Pinel P, van de Moortele P-F, Le Bihan D, Cohen L, Dehaene S (2000) Understanding dissociations in dyscalculia: A brain imaging study of the impact of number size on the cerebral networks for exact and approximate calculation. Brain 123:2240-2255.

Stromswold K, Caplan D, Alpert N, Rauch S (1996) Localization of syntactic comprehension by positron emission tomography. Brain and Language 52:452-473.

Tahmasebi AM, Davis MH, Wild CJ, Rodd JM, Hakyemez H, Abolmaesumi P, Johnsrude IS (2012) Is the link between anatomical structure and function equally strong at all cognitive levels of processing? Cerebral Cortex 22:1593-1603.

Thompson-Schill SL, Bedny M, Goldberg RF (2005) The frontal lobes and the regulation of mental activity. Current Opinion in Neurobiology 15:219-224.

Tomaiuolo F, MacDonald JD, Caramanos Z, Posner G, Chiavaras M, Evans AC, Petrides M (1999) Morphology, morphometry and probability mapping of the pars opercularis of the inferior frontal gyrus: An in vivo MRI analysis. European Journal of Neuroscience 11:3033-3046.

Traxler MJ, Morris RK, Seely RE (2002) Processing subject and object relative clauses: Evidence from eye movements. Journal of Memory and Language 47:69-90.

Trueswell JC, Tanenhaus MK, Garnsey SM (1994) Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. Journal of Memory and Language 33:285-318.

Tyler LK, Marslen-Wilson WD, Randall B, Wright P, Devereux BJ, Zhuang J, Papoutsi M, Stamatakis EA (2011) Left inferior frontal cortex and syntax: Function, structure and behaviour in patients with left hemisphere damage. Brain 134:415-431.

Ullman MT (2004) Contributions of memory circuits to language: The declarative/procedural model. Cognition 92:231-270.

Van Dyke JA, Johns CL (2012) Memory interference as a determinant of language comprehension. Language and Linguistics Compass 6:193-211.

Vandenberghe R, Nobre AC, Price CJ (2002) The response of left temporal cortex to sentences. Journal of Cognitive Neuroscience 14:550-560.

Varley RA, Siegal M (2000) Evidence for cognition without grammar from causal reasoning and "theory of mind" in an agrammatic aphasic patient. Current Biology 10:723-726.

Varley RA, Klessinger NJC, Romanowski CAJ, Siegal M (2005) Agrammatic but numerate. Proceedings of the National Academy of Sciences 102:3519-3524.

Wartenburger I, Heekeren HR, Burchert F, De Bleser R, Villringer A (2003) Grammaticality judgments on sentences with and without movement of phrasal constituents: an event-related fMRI study. Journal of Neurolinguistics 16:301-314.

Wernicke C (1874/1969) The symptom complex of aphasia: A psychological study on an anatomical basis. In: Boston Studies in the Philosophy of Science (Cohen RS, Wartofsky MW, eds), pp 34-97. Dordrecht, The Netherlands: D. Reidel Publishing.

Whorf BL (1965/2012) Language, thought, and reality: Selected writings of Benjamin Lee Whorf. Cambridge, MA: Mit Press.

Willems RM, Benn Y, Hagoort P, Toni I, Varley R (2011) Communicating without a functioning language system: Implications for the role of language in mentalizing. Neuropsychologia 49:3130-3135.

Wilson SM (2014) The impact of vascular factors on language localization in the superior temporal sulcus. Human Brain Mapping 35:4049-4063.

Wilson SM, Saygin AP (2004) Grammaticality judgment in aphasia: Deficits are not specific to syntactic structures, aphasic syndromes, or lesion sites. Journal of Cognitive Neuroscience 16:238-252.

Wilson SM, Molnar-Szakacs I, Iacoboni M (2008) Beyond superior temporal cortex: Intersubject correlations in narrative speech comprehension. Cerebral Cortex 18:230-242.

Wise RJS, Scott SK, Blank SC, Mummery CJ, Murphy K, Warburton EA (2001) Separate neural subsystems within "Wernicke's area". Brain 124:83-95.

Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of Neurophysiology 106:1125-1165.

Zhang LM, Pylkkanen L (2015) The interplay of composition and concept specificity in the left anterior temporal lobe: An MEG study. NeuroImage 111:228-240.

Part A

Language-selective and domain-general mechanisms in language comprehension

[This page intentionally left blank]

Chapter 1

A functional dissociation between language and multiple-demand systems revealed in patterns of BOLD signal fluctuations

With Nancy Kanwisher and Evelina Fedorenko

Abstract

What is the relationship between language and other high-level cognitive functions? Neuroimaging studies have begun to illuminate this question, revealing that some brain regions are quite selectively engaged during language processing whereas other, "multiple-demand (MD)" regions are broadly engaged by diverse cognitive tasks. Nonetheless, the functional dissociation between the language and MD systems remains controversial. Here we tackle this question with a synergistic combination of fMRI methods: we first define candidate language-specific and MD regions in each participant individually (using functional localizers), and then measure BOLD signal fluctuations in these regions during two naturalistic conditions ("rest" and story comprehension). In both conditions, signal fluctuations strongly correlate among language regions, as well as among MD regions, but correlations across systems are weak or negative. Moreover, data-driven clustering analyses based on these inter-region correlations consistently recover two clusters corresponding to the language and MD systems. Thus, although each system forms an internally integrated whole, the two systems dissociate sharply from each other. This independent recruitment of the language and MD systems during cognitive processing is consistent with the hypothesis that these two systems support computationally distinct cognitive functions.

1.1. Introduction

Although the key brain regions engaged in language processing have been known since Broca (1861/2006) and Wernicke (1874/1969), debates continue on whether, and to what extent, they overlap with regions engaged in other cognitive processes. Many neuroimaging studies have reported that brain regions that process language (particularly Broca's area) also engage in many non-linguistic processes, including arithmetic (e.g., Dehaene et al., 1999; Stanescu-Cosson et al., 2000), music perception (Maess et al., 2001; Koelsch et al., 2002), working memory and cognitive control (e.g., Kaan and Swaab, 2002; Koechlin and Jubault, 2006; Hein and Knight, 2008; Blumstein, 2009; January et al., 2009). Yet other studies have found that regions activated during non-linguistic tasks are distinct from language-processing regions (Monti et al., 2009; Fedorenko et al., 2011; Fedorenko et al., 2012a; Monti and Osherson, 2012; Monti et al., 2012). Specifically, many cognitively demanding tasks activate a set of frontal and parietal regions known as the "multiple demand (MD)" system (Duncan and Owen, 2000; Miller and Cohen, 2001; Braver et al., 2003; Cole and Schneider, 2007; Dosenbach et al., 2008; Duncan, 2010; Fedorenko et al., 2013), which does not overlap with the classic fronto-temporal language system. Nonetheless, the dissociation between a putatively language-specific system and this domain-general MD system remains controversial (Thompson-Schill et al., 2005; Blumstein and Amso, 2013).

To test for this dissociation, here we compared the BOLD signal time-courses of candidate language and MD regions by synergistically combining two fMRI methods: functional localizers and functional correlations. First, we functionally localized candidate regions of interest in each participant (Fedorenko et al., 2010; Fedorenko et al., 2013) using tasks that target linguistic processing (language localizer) and cognitive effort (MD localizer). Next, participants were scanned during a "rest" period or during a story comprehension task (conditions that are independent of, and less constrained than, the localizers). The time-course of BOLD signal fluctuations during these two conditions was then extracted from each functionally-defined region. Finally, we measured the pairwise correlations between time-courses of different regions (separately for each condition). This approach enabled us to answer three questions: (i) To what extent do candidate language regions form a functionally integrated system (Cordes et al., 2000; Hampson et al., 2002; Turken and Dronkers, 2011; Newman et al., 2013; Yue et al., 2013), as indexed by high correlations among these regions?; (ii) To what extent do candidate MD regions similarly form an integrated system? (Dosenbach et al., 2007; Seeley et al., 2007; Hampshire et al., 2012); and critically, (iii) How functionally dissociable are language and MD regions from each other, as indexed by weak (or negative) correlations between pairs of regions straddling the two systems?

The current approach harnesses the complementary strengths of functional localizers and functional correlations. First, the "rest" and story comprehension conditions allow us to sample a broader, more naturalistic range of cognitive processes compared to task-based studies. Second, functional correlations allow us to use not only hypothesis-driven methods but also data-driven clustering to discover the relationship between language and MD regions based on the co-variation of their respective signal time-courses. Third, we can straightforwardly interpret the emerging clusters in terms of their functionally characterized constituents, because our regions of interest are functionally localized. Therefore, we do not have to rely on "reverse inference" from stereotaxic coordinates (Poldrack, 2006), which is inevitable when no functional localizers are used (Yeo et al., 2011; Lee et al., 2012; Mantini et al., 2013; Tie et al., 2014).

1.2. Materials and methods

1.2.1. Participants

Eighteen adult participants (6 males) aged 18-30 participated in a resting-state scan. Ten participants (3 males) aged 18-30 were scanned during a story comprehension task (six of these also participated in the resting-state scan). All 22 participants also completed independent localizer runs that were used to define candidate language and MD regions. Participants were right-handed, native speakers of English from MIT and the surrounding Cambridge community. All provided informed consent in accordance with the requirement of MIT's Committee on the Use of Humans as Experimental Subjects (COUHES).

1.2.2. Functional localization of candidate language and MD regions

1.2.2.1. Data acquisition

Structural and functional data were collected on a whole-body 3 Tesla Siemens Trio scanner with a 32-channel head coil at the Athinoula A. Martinos Imaging Center at the McGovern Institute for Brain Research at MIT. T1-weighted structural images were collected in 176 sagittal slices (1mm isotropic voxels; TR = 2530ms; TE = 3.48ms). Functional blood oxygenation level dependent (BOLD) data were acquired using an EPI sequence with a flip angle of 90° and applying GeneRalized Autocalibrating Partially Parallel Acquisition (GRAPPA) with an acceleration factor of 2. Images were collected in 31 near-axial slices, acquired in an interleaved order with a

10% distance factor (in-plane resolution: 2.1×2.1mm; slice thickness: 4mm; FoV: 200mm in the phase encoding (A>>P) direction; matrix size: 96mm×96mm; TR: 2000ms; TE: 30ms). Prospective Acquisition CorrEction (PACE; Thesen et al., 2000) was used to adjust the positions of the gradients based on the participant's head motion one TR back. The first ten seconds of each run were excluded to allow for steady state magnetization.

1.2.2.2. Design, materials and procedure

Candidate language and MD regions of interest were functionally defined. The language localizer has been previously demonstrated to identify language-selective brain regions (Fedorenko et al., 2011) sensitive to high-level linguistic information, including syntax and lexical semantics (Fedorenko et al., 2012b) and, accordingly, contrasted reading of sentences with reading of sequences of pronounceable nonwords (presented one word/nonword at a time). After each sentence/nonword sequence, a probe word/nonword appeared and participants had to decide whether the probe item appeared in the preceding stimulus. Each of our 22 participants completed 2-4 runs of the localizer, with sentence and nonwords blocks lasting either 24s (4 blocks per condition per run in a 336s run, 4 participants) or 18s (8 blocks per condition per run in a 378s run, 12 participants; or 6 blocks per condition per run in a 396s run, 6 participants). The order of sentence and nonwords blocks was counterbalanced across runs and participants. This localizer (available from http://web.mit.edu/evelina9/www/funcloc/funcloc_localizers.html) is robust to changes in materials, task and modality of presentation (Fedorenko et al., 2010).

For the MD localizer, participants performed a spatial working memory task that we have found to broadly and robustly activate the MD system (Fedorenko et al., 2013). Participants had to keep track of four (easy condition) or eight (hard condition) locations in a 3×4 grid (Fedorenko et al., 2011). In both conditions, participants performed a two-alternative forced-choice task at the end of each trial to indicate the set of locations they just saw. The contrast *hard* > *easy* targets brain regions engaged in cognitively demanding tasks; Fedorenko et al. (2013) have shown that the regions activated by this task are also activated by a wide range of other tasks contrasting a difficult versus an easy condition. Each of our 22 participants completed 1-3 runs of this MD localizer, with hard and easy blocks lasting either 34s (5 blocks per condition per run in a 436s run, 10 participants), 32s (6 blocks per condition per run in a 448s run, 11 participants) or 18s (6 blocks per condition per run in a 288s run, 1 participant). The order of hard and easy blocks was counterbalanced across runs and participants.

1.2.2.3. Data preprocessing

MRI data were spatially preprocessed using SPM5 (http://www.fil.ion.ucl.ac.uk /spm). To reduce data transformations, we performed all analyses in native functional space. Each participant's functional data were motion-corrected and then smoothed with a 4mm FWHM Gaussian filter. In addition, the anatomical image of each participant was segmented into three probability maps denoting areas of gray matter, white matter and cerebro-spinal fluid (CSF) (see section 1.2.3.3), and these maps were then co-registered to the native functional space. The anatomical images were also further used for cortical surface reconstruction (see next section).

1.2.2.4. Definition of group-constrained, participant-specific fROIs

Data for the localizer tasks were modeled using a general linear model with a boxcar regressor modeling entire blocks convolved with a canonical HRF function. For each participant, functional regions of interest (fROIs) were defined by combining the activation map for the localizer contrast with group-level constrains ("masks") (Fedorenko et al., 2010; Julian et al., 2012). The masks demarcated brain areas within which most or all individuals in prior studies showed activity for the localizer contrasts (see Table 1.1).

For the language fROIs, we used masks derived from a group-level representation of data for the *sentences > nonwords* contrast in an independent group of subjects (Fedorenko et al., 2010) (masks available for download at http://web.mit.edu/evelina9/www/funcloc/funcloc _parcels.html). Following Fedorenko et al. (2011), eight masks were used for defining fROIs in the left hemisphere, including regions in the posterior and middle-posterior temporal lobe (PostTemp and MidPostTemp, respectively), anterior and middle-anterior temporal lobe (AntTemp and MidAntTemp, respectively), inferior and orbital-inferior frontal gyrus (IFG and IFGorb, respectively), middle frontal gyrus (MFG) and angular gyrus (AngG). These masks were mirror-projected onto the right hemisphere to create eight homologous masks, resulting in a total of 16 language masks. As the masks (denoted by gray contours in Figure 1.1) cover significant parts of the cortex, their mirrored version is likely to encompass the right hemisphere homologues of the left-hemispheric language-selective regions, despite possible hemispheric asymmetries in their precise anatomical location.

The right-hemisphere homologues were included for two reasons. First, these regions appear to be activated during at least some aspects of language processing, albeit usually not as strongly as the typical, left-lateralized language regions (Chiarello et al., 2003; Jung-Beeman, 2005). Second, given that left hemisphere damage, but typically not right hemisphere damage, leads to difficulties in language production and comprehension (Geschwind, 1970; Damasio, 1992), we wanted to examine hemispheric differences in functional correlations. For example, we wanted to see whether left hemisphere language regions would show stronger inter-region correlations, which might be expected of a core language system (cf. the MD regions, which are strongly bilateral and should thus show similarly strong inter-region correlations in the left and right hemispheres).

For the MD fROIs, we used anatomical masks (Tzourio-Mazoyer et al., 2002) that included the main regions linked to MD activity in prior work, following the methods of Fedorenko et al. (2013) (for a similar approach, see Fedorenko et al., 2012a). Eighteen masks were used, 9 in each hemisphere, including regions in the opercular inferior frontal gyrus (IFGop), middle frontal gyrus including its orbital part (MFG and MFGorb, respectively), insular cortex (Insula), precentral gyrus (PrecG), supplementary and pre-supplementary motor area (SMA), inferior and superior parietal lobe (ParInf and ParSup, respectively) and anterior cingulate cortex (ACC).

These group-level masks, in the form of binary maps, were used to systematically constrain the selection of participant-specific fROIs. Thus, for each participant, 16 candidate language fROIs were created by intersecting the participant's unthresholded t-map for the *sentences* > *nonwords* contrast with the language masks. For each mask, the 10% of voxels with the highest t-values in the intersection image were then used to define a fROI (note that the voxels included in the right hemisphere fROIs were not constrained to be mirrored versions of their left hemisphere counterparts, but were only constrained to land within a mirrored version of the broad masks). Similarly, 18 candidate MD fROIs were created for each participant by intersecting the participant's unthresholded t-map for the *hard* > *easy* contrast with the MD masks, again selecting the 10% of voxels with the highest t-values within each mask. Finally, we excluded a small set of voxels that were contained in more than one fROI due to small spatial overlap between language and MD activation maps. Across participants, these excluded voxels comprised 1.67%(2.60) of our fROIs (for all participants, 16 fROIs had no voxels excluded from them; each of the remaining 18 fROIs had less than 7% excluded voxels).

Defining fROIs by choosing the 10% of voxels with the highest t-values for a localizer contrast in a given mask balances the trade-off between: (i) choosing only voxels with a BOLD time-course that strongly co-varies with the localizer conditions (sentences vs. non-words or hard vs. easy); and (ii) having a sufficient number of voxels in the fROI. In addition, this procedure ensures that each fROI has a constant size across participants. However, we obtained similar results to those reported below when fROIs were instead defined by intersecting the language or

MD masks with a thresholded *t*-map for the language or MD contrasts, respectively (p=0.001, uncorrected).

The language and MD masks used here were originally created in Montreal Neurological Institute (MNI) space. Therefore, prior to defining fROIs, the masks had to be projected onto each participant's native functional space. This was done in two steps: first, Combined Volume and Surface (CVS) registration (Postelnicu et al., 2009) was used to estimate the transformation of an MNI anatomical (T1) template to the native anatomical space of each subject, and the resulting transformation was applied to the masks. Second, affine co-registration was used to project the masks from native anatomical space onto native functional space. Only then were the masks intersected with participants' *t*-maps from the functional localizers.

The localizer effects were highly reliable in all fROIs. Reliability was tested via an *n*-fold, leave-one-out cross validation across runs: for each participant, we defined fROIs based on all localizer runs but one, and then derived estimates of the localizer contrast effect for the left-out run in these fROIs. The contrast effect estimates were averaged across all possible left-out runs and then tested for significance across participants (FDR-corrected for the number of regions;)(FDR-corrected for the number of regions; Benjamini and Yekutieli, 2001). For all left-hemisphere candidate language regions, $t_{(21)}>6.47$, $p<10^{-5}$. For all right-hemisphere candidate language regions, $t_{(21)}>2.70$, $p<0.007$. For all candidate MD fROIs, $t_{(19)}>5.64$, $p<10^{-4}$. Figure 1.1 shows the language and MD fROIs in the left hemisphere of 3 representative subjects, as well as probability maps of fROI locations across all subjects.

Table 1.1. Functional regions of interest (fROIs)

Candidate Language Regions				**Candidate MD Regions**			
Left Hemisphere		Right Hemisphere		Left Hemisphere		Right Hemisphere	
1	MidPostTemp	9	MidPostTemp	17	IFGop	26	IFGop
2	PostTemp	10	PostTemp	18	MFG	27	MFG
3	MidAntTemp	11	MidAntTemp	19	MFGOrb	28	MFGOrb
4	AntTemp	12	AntTemp	20	ParInf	29	ParInf
5	IFG	13	IFG	21	ParSup	30	ParSup
6	IFGOrb	14	IFGOrb	22	ACC	31	ACC
7	MFG	15	MFG	23	Insula	32	Insula
8	AngG	16	AngG	24	SMA	33	SMA
				25	PrecG	34	PrecG

MD, multiple-demand; MidPostTemp, middle-posterior temporal lobe; PostTemp, posterior temporal lobe; MidAntTemp, middle-anterior temporal lobe; AntTemp, anterior temporal lobe; IFG, inferior frontal gyrus; IFGOrb, orbital-IFG; MFG, middle frontal gyrus; AngG, angular gyrus; IFGop, opercular IFG; MFGOrb, orbital-MFG; ParInf, inferior parietal lobe; ParSup, superior parietal lobe; ACC, anterior cingulate cortex; Insula, insular cortex; SMA, supplementary and presupplementary motor area; PrecG, precentral gyrus.

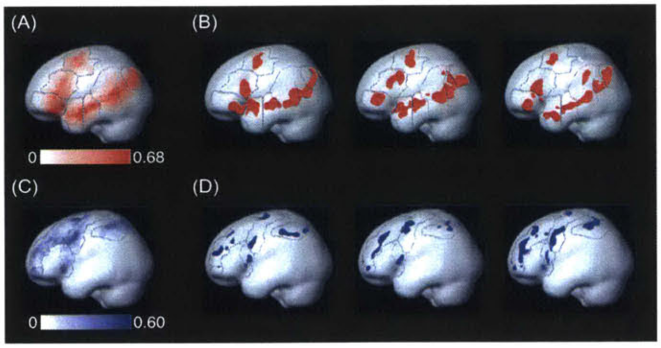

Figure 1.1. Group-constrained, participant-specific fROIs. (A) and (C) show probability maps of the locations of fROIs across participants, for the language (red) and MD (blue) systems, respectively. Higher color saturation corresponds to a higher number of participants having a significant activation in the relevant voxel. Apparent overlap between fROIs is only at the group level, not the individual participant level. (B) and (D) show candidate language fROIs (red) and candidate MD fROIs (blue) in the left hemisphere of 3 representative participants. In all subfigures, dark grey lines demarcate the masks used to constrain the location of fROIs (see section 1.2.2.4). Images are in MNI space for illustration purposes only (fROI definition and functional correlation analyses were carried out in the native functional space of each individual participant). Overlap between different fROIs is only apparent due to the projection of fROIs onto the cortical surface.

1.2.3. Functional correlation analysis

1.2.3.1. Data acquisition

Functional data were collected using the same parameters as for the functional localizers.

1.2.3.2. Design, materials and procedure

In the resting state condition, participants were instructed to close their eyes but to remain awake and let their mind wonder for 5mins. In the story comprehension condition, participants listened to 4-6 stories over the scanner-safe headphones (Sensimetrics, Malden, MA). Each story lasted between 4.5 and 6mins. Stories were constructed from existing publicly available texts (fairy-tales, short stories and Wikipedia articles) but edited so as to include a variety of linguistically interesting phenomena that do not occur with sufficiently high frequency in natural texts (e.g., infrequent words, non-local syntactic dependencies, unusual syntactic constructions, temporary ambiguity, etc.; for examples, see Table 1.2). (The motivation for editing the stories had to do with the use of these materials in another project aimed at understanding the processing of different kinds of linguistic complexity.) The stories were then recorded by two native English speakers (one male, one female). After each story, participants answered six comprehension questions, presented in a two-alternative forced-choice format. For each participant, accuracy on these questions was significantly above chance, as indicated by the binomial test (for all tests, $p<10^{-11}$, Bonferroni corrected for the number of participants; mean accuracy across participants: 83.09%(10)).

1.2.3.3. Data preprocessing

Functional data were spatially preprocessed using the same procedure applied for the localizer runs. Temporal preprocessing was carried out using the CONN toolbox (Whitfield-Gabrieli and Nieto-Castañón, 2012) with default parameters unless otherwise specified. First, noise introduced by signal fluctuations originating from non-neuronal sources (e.g., respiratory and cardiac activity) was removed. To this end, the first 5 temporal principal components of the BOLD signal time-course extracted from the white matter were regressed out of each voxel's time-course; signal originating in the CSF was similarly regressed out. White matter and CSF voxels were identified based on segmentation of the anatomical image (Behzadi et al., 2007). The first 6 principal components of the six motion parameters estimated during offline motion correction were also regressed out, as well as their first time-derivative. Second, the residual signal was bandpass filtered (0.008-0.09Hz) in order to only preserve low-frequency signal fluctuations (Cordes et al., 2001).

Table 1.2. Linguistic materials used in the story comprehension task.

	Examples
Infrequent words	"Autosomal"
	"Brunt"
	"Conjectured"
Non-local syntactic dependencies	"The kindly Lord of the Manor who the people had often asked for help..."
	"The severity of the problem the people faced..."
	"The water snail that she had discovered a couple of days ago..."
Unusual syntactic constructions	"A source of great trouble to the local folk the boar was..."
	"It was the first huntsman who was..."
	"Into vapor the water drops that danced in the ocean had been changed..."
Temporary ambiguity	"The huntsman questioned by the Lord..."
	"The matron understood my idea was something that I was excited about..."
	"Abby's mom denied Abby's version of the story was true..."

1.2.3.4. Data analysis

Analysis of functional correlations was carried out separately for each of the two experimental conditions (resting state and story comprehension). For each participant, we averaged the BOLD signal time-course across all voxels in each fROI. For each pair of fROIs, Pearson's moment correlation coefficient was then computed between their respective time-courses. These correlations were Fisher-transformed to improve normality, and three one-way, repeated-measures analyses of variance (ANOVAs) were then performed on the data in order to assess whether different regions showed different patterns of functional correlations. First, we compared the average correlation within the candidate language system (i.e., the average of all 120 pairwise correlations among the 16 language fROIs), the average correlation within the candidate MD system (i.e., the average of all 153 pairwise correlations among the 18 MD fROIs), and the average correlation across the two systems (i.e., the average of all 288 pairwise correlations between a language fROI and a MD fROI). Second, for the language system, we compared the mean correlation within the left-hemisphere (averaging across all 28 pairwise correlations among the 8 language fROIs in the left-hemisphere), the analogous mean correlation within the right-hemisphere, and the mean correlation across the two hemispheres (averaging across 64 pairwise correlations). Third, the same within-hemisphere vs. across-hemispheres comparison was carried out for the MD system.

For visualization purposes, we also created two group-level matrices of fROI-to-fROI correlations, one for each experimental condition. Specifically, the Fisher-transformed correlation between each pair of fROIs was averaged across participants (the Fisher-transform decreases the bias in averaging; Silver and Dunlap, 1987), and the resulting average correlations were then inverse-Fisher-transformed. The two group-level fROI-to-fROI correlation matrices are presented in Figure 1.2. However, to ensure that the patterns of functional correlations reported here were consistently observed across individual participants, the majority of our analyses did not use these average correlation matrices but were instead performed within participants (e.g., the repeated-measures ANOVAs described above). Only the hierarchical clustering analysis (see below) relied on group-level average correlations.

1.2.3.5. Controlling for the effects of head motion

Head motion artifacts have been reported to affect functional correlations (Van Dijk et al., 2010; Power et al., 2012; Satterthwaite et al., 2012). To ensure that such artifacts could not account for the results reported here, we performed two control analyses. First, for each experimental run of each participant, time-points with excessive head motion ("motion spikes") were identified using the Artifact Detection Toolbox implemented in MATLAB (available for download at http://www.nitrc.org/projects/artifact_detect/). Each motion spike was then included as a regressor during temporal preprocessing (see section 1.2.3.3), thus removing the effects of these time-points on the BOLD time courses (Lemieux et al., 2007; Satterthwaite et al., 2013). This control analysis and our original analysis (without "spike regression") resulted in qualitatively similar patterns of functional correlations. To minimize data manipulations, we report the results of our original analysis.

Second, we tested whether individual differences in estimated head motion could explain individual differences in functional correlation patterns. First, each participant's six motion parameters were collapsed to a single value, Mean Relative Displacement (MRD; Jenkinson et al., 2002). Then, we computed the Pearson correlation, across participants, between MRD and each of three measures described above: (i) the mean functional correlation within the language system; (ii) the mean functional correlation within the MD system; and (iii) the mean functional correlation across the two systems. None of these measures were significantly correlated with MRD (FDR-corrected) in either the resting state or the story comprehension conditions (the correlations that were significant prior to FDR correction were opposite in direction to our reported effects). On average, individual differences in head motion explained 5% of the individual differences in functional correlation patterns. We then repeated this analysis using the point bi-serial correlation instead of the Pearson correlation by splitting our sample into a "high MRD" half and a "low MRD" half. While the two halves significantly differed in MRD, they did not differ in their functional correlation patterns. We therefore conclude that our results reported below cannot result from a head motion artifact.

1.2.4. Clustering analyses

1.2.4.1. K-means

In order to reveal the dominant patterns of functional correlations across our fROIs in a relatively data-driven fashion, we submitted the average BOLD signal time-courses from candidate language and MD fROIs to Matlab's *K*-means clustering algorithm, separately for each participant. To ensure the choice of *K* did not impose on the data an implicitly hypothesized division into language and MD systems, separate analyses were run with values of *K* ranging from 2 to 6 clusters. The pairwise distance measure used for clustering was defined as 1 minus the pairwise correlation between time-courses. To choose the initial cluster centroids, *K* time-courses out of a participant's data were randomly sampled, and this procedure was repeated 50 times to generate multiple clustering solutions. To pool the resulting data we then computed, for each pair of fROIs, the probability (percentage of solutions across random initializations and participants) that the two fROIs would both be assigned to the same cluster. Such pooling provides a straightforward way to collapse results across participants with different clustering solutions or cluster-numbering (i.e., order) (such pooling is also known as "consensus clustering"; see Lancichinetti and Fortunato, 2012; Bassett et al., 2013).

We assessed the significance of our clustering results with a permutation test. For each participant, we created a surrogate BOLD signal time-course for each fROI by phase-shuffling its original time-course (i.e., reassigning the phases of different frequencies uniformly at random, with replacement). The 34 surrogate time-courses of each subject were then clustered and the clustering solutions were pooled across participants, using the same procedures described above. We repeated this clustering of surrogate time-courses 1,000 times, generating for each pair of fROIs a null distribution of the probability that they would be assigned to the same cluster. The true probability, based on clustering the real data, was compared against this distribution to produce a two-tailed *p*-value. Multiple comparisons were FDR-corrected separately for each fROI (each pair of the 34 fROIs was assigned a *p*-value, so each fROI had 33 *p*-values associated with it). This test yielded results similar to those obtained when we shuffled the original time-courses across fROIs instead of generating surrogate time-courses.

In the previous analysis, the clustering algorithm was data-driven in the sense that it was not provided with information about which fROIs were candidate language regions and which were candidate MD regions. Nonetheless, the analysis was still constrained to treat each fROI as a distinct entity, as the clustering was run on time-courses that were averaged across all voxels within each fROI. To partially relax this constraint, our next analysis clustered the BOLD signal time-courses of individual voxels across all fROIs. As in the previous analysis, 50 clustering solutions were generated for each participant, where randomly sampled time-courses served as initial cluster centroids. For each clustering solution, we then performed the following computation: first, we examined each set of voxels originating from within a single fROI to determine its “dominant cluster” (that is, the cluster that had the largest number of voxels in that fROI assigned to it). Then, for each pair of fROIs, we computed the percentage of voxels in the first fROI assigned to the cluster that was dominant in the second fROI (this resulted in two measures, depending on which fROI was "first" and which was "second"). This procedure provided, for each pair of fROIs, a voxel-wise measure of cluster similarity, which was then averaged across random initializations and participants. To test the significance of this cluster similarity measure for each pair of fROIs, we applied a phase-shuffling permutation test following the same procedure as described above.

This second analysis used data from single voxels rather than fROIs, yet the similarity measure we computed for pooling the results across participants still referred to the original grouping of voxels into fROIs. However, this measure was only computed after voxel-wise clustering had taken place. We chose this measure because our data did not allow for comparing single voxels across participants: as fROIs were defined in a participant-specific manner, voxels falling within a fROI for one participant might not have fallen within any fROI for another participant, thus entering the clustering analysis for the former participant but not the latter.

1.2.4.2. *Hierarchical clustering*

Hierarchical clustering is an algorithm that creates a binary tree-structure connecting elements in a set, such that the length of branches on the tree approximates the distances among the elements as provided by the user (Hartigan, 1975). Clustering together elements whose connecting path on the tree is shorter than a chosen length therefore creates a partition of the element set without pre-specifying the number of resulting clusters (in contrast to *K*-means). We performed hierarchical clustering on our fROIs, providing the group-level fROI-to-fROI correlation matrix as input, so that the distance between two fROIs was defined as 1 minus their correlation. Clustering was based on average linkage, such so that two clusters were merged into a bigger cluster based on the mean distance between their respective members.

The optimal partition of fROIs based on the resulting tree was identified via a measure of modularity (Newman and Girvan, 2004). First, by gradually decreasing the path-length used as a criterion for clustering fROIs, we generated the set of all possible partitions licensed by our hierarchical clustering solution (the longest path-length generates a single cluster consisting of all 34 fROIs; the shortest path-length generates 34 singleton clusters). Then, for each partition, we computed a modularity measure that is appropriate for detecting clusters in correlated data (Gómez et al., 2009). High modularity values indicate clustering solutions where, within each cluster, the positive functional correlations are stronger (and the negative functional correlations are weaker) compared to what is expected under a null model. The null model is a random fROI-to-fROI correlation matrix that preserves, for each fROI, the sum of its positive correlations and the sum of its negative correlations with the other fROIs.

1.3. Results

1.3.1. Functional correlation analysis: comparing systems and hemispheres

1.3.1.1. Comparison of the language and MD systems

Figure 1.2 shows, for the resting state and story comprehension conditions, matrices of pairwise correlations between candidate language and MD fROIs computed on the time-courses of BOLD signal fluctuations. A clear partition of the fROIs is visually evident prior to any statistical analysis: most pairs of language fROIs are strongly and positively correlated with each other (cf. top-left quadrant of the correlation matrices), and most pairs of MD ROIs are also strongly and positively correlated with each other (cf. bottom-right quadrant of the correlation matrices); but correlations of most pairs consisting of a language fROI and a MD fROI are noticeably weaker (cf. bottom-left and top-right quadrants of the correlation matrices). Moreover, during story comprehension there are significant negative correlations between candidate left-hemisphere language fROIs and right-hemisphere MD fROIs. These results indicate a functional architecture comprised of two functional networks, one consisting of language regions and the other consisting of MD regions. BOLD signal fluctuations within each network are highly synchronized, but the two networks are functionally dissociated.

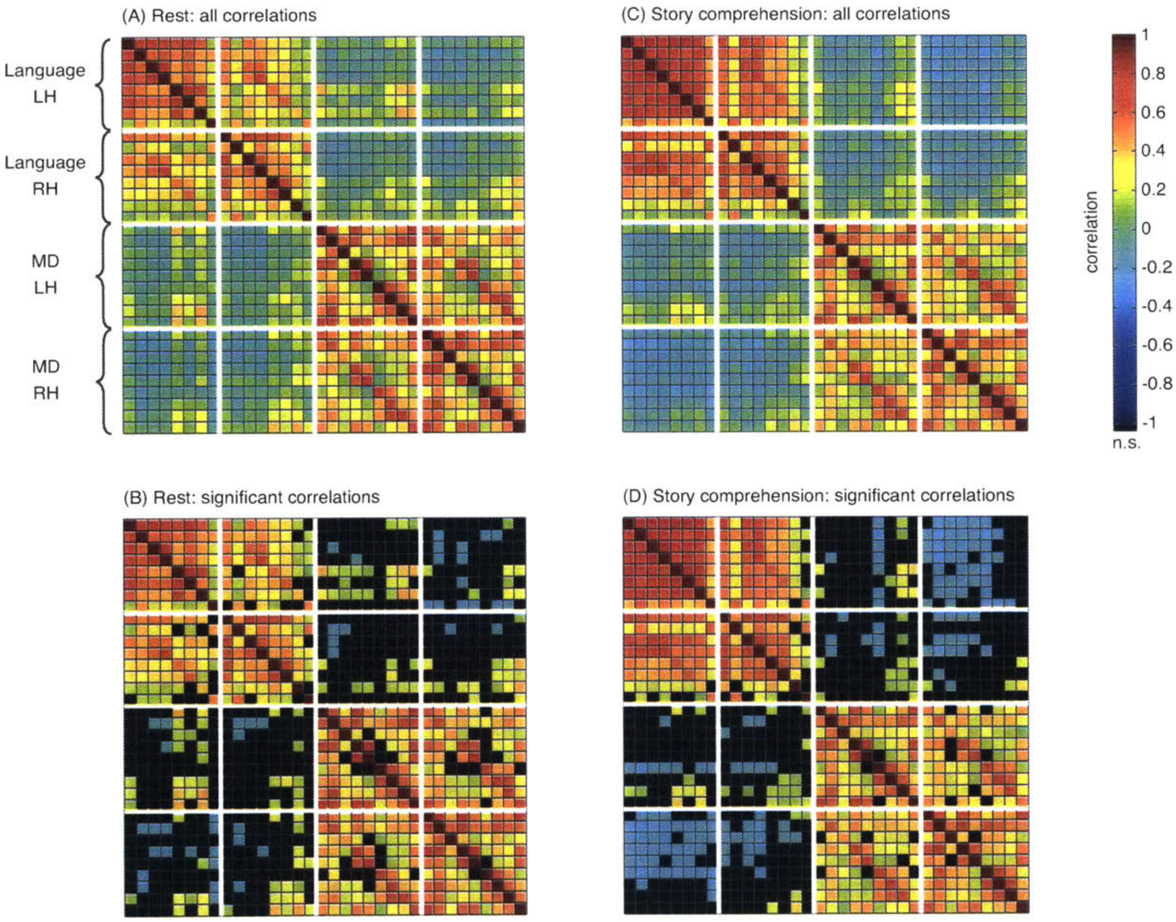

Figure 1.2. Matrices of fROI-to-fROI functional correlations, for (A) the resting state condition; and (C) the story comprehension condition. Matrices in (C) and (D) present the same data as (A) and (B) respectively, but show only significant correlations (α=0.05, FDR-corrected). Non-significant correlations are colored in black. The order of fROIs across rows (and columns) follows Table 1.1, where regions are sorted by system (language, then MD). Within each system, fROIs are sorted by hemisphere (left hemisphere (LH), then right hemisphere (RH)). Thick white lines separate these subsets of fROIs.

To quantitatively test for the language-MD functional dissociation, we compared the average pairwise correlation within the language system (across all fROI pairs) to the average pairwise correlation within the MD system and the average pairwise correlation across the two systems (Figure 1.3) (averages were computed based on Fisher-transformed correlations; see Materials and Methods). Specifically, a one-way, repeated-measures ANOVA was carried out to compare functional correlations among these three levels ("within language", "within MD" and "across systems"). Consistent with our qualitative observations, a highly robust effect was revealed in both conditions (resting state: $F_{(2,34)}=62.84$, $p<10^{-11}$; story comprehension: $F_{(2,18)}=78.56$, $p<10^{-8}$). In the resting state condition, post-hoc pairwise comparisons (Bonferroni corrected, for all that follow) showed that the average correlation across the two systems ($r=0.03$, $SD=0.11$ across participants) was weaker than the average correlation within the language system ($r=0.38$, $SD=0.13$) and within the MD system ($r=0.41$, SD=0.09) (for both tests, $t_{(17)}>8.47$, $p<10^{-6}$). Similarly, in the story comprehension condition, the average correlation across the two systems ($r=-0.03$, $SD=0.10$) was weaker than the average correlation within the MD system ($r=0.37$, $SD=0.11$), which was in turn weaker than the average correlation within the language system ($r=0.49$, $SD=0.08$) (for both tests, $t_{(9)}>3.68$, $p<0.016$).

(a) Rest

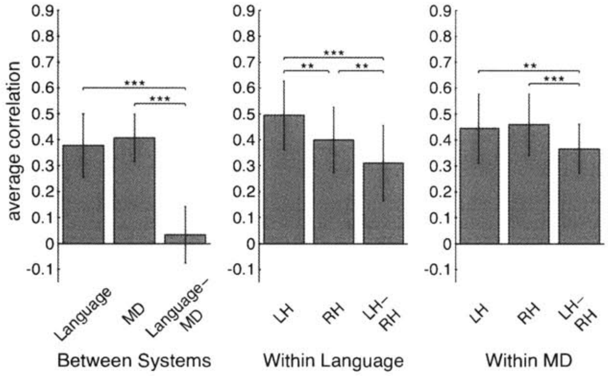

(b) Story Coprehension

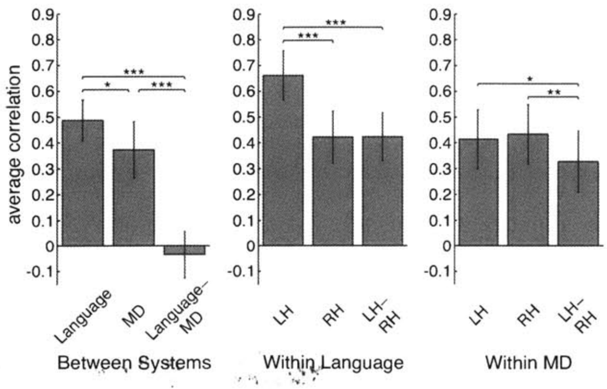

Figure 1.3. Comparisons of average correlations within and across systems and hemispheres, for (A) the resting state condition; and (B) the story comprehension condition. Three repeated-measures comparisons are presented. Left: comparing the average pairwise correlation within the language system (i.e., across all language fROI pairs), the average correlation within the MD system (i.e., across all MD fROI pairs), and the average correlation between the two systems (i.e., across all pairs of a language fROI and a MD fROI). Middle: comparing the average pairwise correlation within the left hemisphere (LH), within the right hemisphere (RH) and between hemispheres, in the language system. Right: comparing the average pairwise correlation within the left hemisphere, within the right hemisphere and between hemispheres, in the MD system. Errorbars show standard deviations across subjects. *$p<0.05$, **$p<0.01$, ***$p<0.001$ (Bonferonni-corrected for multiple comparisons).

Within-system and between-system correlations were also compared while controlling for the effect of anatomical distance among fROIs. To this end, we computed the Euclidean distances between each pair of ipsilateral fROIs, based on either their respective center-of-mass coordinates or their point of maximal proximity. Then, for each fROI, its distances from all other ipsilateral fROIs were regressed out from the corresponding functional correlations (separate analyses were performed for the two distance measures). For each fROI, we then tested whether its residual correlations with other fROIs that belonged to its own system were stronger than its residual correlations with fROIs that belonged to the other system. We found that correlations within each system remained stronger than correlations across the two systems. This difference in correlation strength reached significance for all fROIs (Bonferroni-corrected for multiple comparisons) except for two language fROIs in the right hemisphere: the AngG (in both the resting state and story comprehension conditions) and the MFG (in the resting state condition).

1.3.1.2. Comparison of the left and right hemispheres

We next compared, for each of the language and MD systems, the average pairwise correlation within the left hemisphere to the average pairwise correlation within the right hemisphere and the average pairwise correlation across the two hemispheres (Figure 1.3). Specifically, a one-way, repeated-measures ANOVA was carried out to compare functional correlations among these three levels ("within right hemisphere", "within left hemisphere" and "across hemispheres").

For the language system, a significant hemispheric difference was revealed in both the resting state ($F_{(2,34)}=28.97$, $p<10^{-7}$) and the story comprehension ($F_{(1.11,\ 9.99)}=63.95$, $p<10^{-5}$, Greenhouse-Geisser corrected for non-sphericity) conditions. In the resting state condition, post-hoc pairwise comparisons showed that the average correlation in the left-hemisphere ($r=0.50$, $SD=0.14$) was stronger than the average correlation in the right-hemisphere ($r=0.40$, $SD=0.13$), which was in turn higher than the average correlation across hemispheres ($r=0.31$, $SD=0.15$) (for all tests, $t_{(17)}>3.61$, $p<0.007$). In the story comprehension condition, post-hoc pairwise comparisons showed that the average correlation in the left-hemisphere ($r=0.66$, $SD=0.10$) was stronger than the average correlation in the right-hemisphere ($r=0.42$, $SD=0.11$) as well as across hemispheres ($r=0.43$, $SD=0.10$) (for both tests, $t_{(9)}>7.54$, $p<0.001$), but the latter two did not significantly differ.

In the MD system, significant hemispheric effects were also revealed in the resting state ($F_{(1.22,\ 20.69)}=8.19$, $p=0.007$, Greenhouse-Geisser corrected) and story comprehension ($F_{(2,18)}=13.51$, $p<0.001$) conditions. In the resting state condition, post-hoc pairwise comparisons showed that the average correlation across hemispheres ($r=0.37$, $SD=0.10$) was weaker than the average correlation within the left hemisphere (r=0.44, SD=0.14) and within the right hemisphere (r=0.46, SD=0.12) (for both tests, $t_{(17)}$>4.27, p<0.002). Similar results were found for the story comprehension condition (across hemispheres: r=0.33, SD=0.13; left hemisphere: r=0.42, SD=0.12; right hemisphere: r=0.43, SD=0.12; for both tests, $t_{(9)}$>3.88, p<0.012).

1.3.2. *K*-means

1.3.2.1. Clustering fROIs

We clustered fROIs based on the correlations among their respective average BOLD signal time-courses, separately for each participant. For both the resting state and the story comprehension conditions, clustering the fROIs into K=2 clusters revealed a clear partition between the language and the MD systems (Figure 1.4A). Across participants in the resting state condition, an average of 14.24(1.44) fROIs out of the 16 candidate language fROIs, or 89%(0.09), were grouped into one cluster; while an average of 16.39(1.32) fROIs out of the 18 candidate MD fROIs, or 91%(0.07), were grouped into a different cluster. Similarly, across participants in the story comprehension condition, an average of 14.59(1.00) fROIs out of the 16 candidate language fROIs, or 91%(0.06), were grouped into one cluster; while an average of 16.49(1.80) fROIs out of the 18 candidate MD fROIs, or 92%(0.1), were grouped into a different cluster. A notably inconclusive clustering pattern was only observed for two candidate, right-hemisphere homologue language fROIs, namely the right AngG and MFG. Across participants in the resting state condition, these two regions were assigned to the language-dominant cluster only on 57.2%(47) and 51.8%(48) of the clustering solutions, respectively. Across participants in the story comprehension condition, these two regions were assigned to the language-dominant cluster only on 63.5%(28) and 50.5%(34) of the clustering solutions, respectively. Importantly, the separation between language and MD systems did not result from constraining the algorithm to generate two clusters; a similar pattern was obtained for solutions ranging from 3 to 6 clusters, where candidate language fROIs were still clustered with each other more often than with MD fROIs, and vice versa (Figure 1.5).

The partition of fROIs into a language cluster and a MD cluster was not expected under an empirical null model based on surrogate BOLD time-courses (created via phase-shuffling the original data). Out of the 91 possible pairs of 14 language fROIs (excluding the right AngG and MFG), 86 pairs (94.5%) in the resting state condition and 91 pairs (100%) in the story comprehension condition were jointly clustered more often than expected by chance. Similarly, out of the 153 possible pairs of 18 MD fROIs, 137 pairs (89.5%) in the resting state condition and

153 pairs (100%) in the story comprehension condition were jointly clustered more often than expected by chance. Conversely, out of 288 possible pairs consisting of a language fROI and a MD fROI, 285 pairs (99%) in the resting state condition and 288 pairs (100%) in the story comprehension condition were jointly clustered less often than expected by chance.

1.3.2.2. Clustering individual voxels

When BOLD time-courses from all individual voxels within the fROIs were clustered into K=2 clusters, a "language-dominant" cluster and a "MD-dominant" cluster again emerged. A high percentage of voxels originating within language fROIs were all assigned to one cluster (resting state: 76.59%(9.11); story comprehension: 83.76%(6.44); averaged across participants); but this same cluster consisted of much fewer voxels originating within MD fROIs (resting state: 19.66%(7.67); story comprehension: 19.38%(9.14)). This cluster was hence language-dominant, with the other cluster showing the opposite, MD-dominant pattern. The difference between the percentage of language voxels vs. MD voxels assigned to the same cluster was significant (resting state: $t_{(17)}=16.61, p<10^{-11}$; story comprehension: $t_{(9)}=15.28, p<10^{-7}$) (Figure 1.4C).

When we compared, for every given pair of fROIs, the assignments of their constituent voxels into the two clusters, we observed a clustering pattern similar to that reported for fROI-wise clustering. Namely, in both the resting state and the story comprehension conditions, a partition emerged between the language and MD systems (Figure 1.4B). Thus, the majority of voxels originating in language fROIs were all assigned to the same cluster, whereas the majority of voxels originating in MD fROIs were assigned to the other cluster. The least conclusive clustering pattern was again observed for the right AngG and MFG (candidate right-homologue language fROIs). Only about half of the voxels originating in these fROIs were assigned to the language-dominant cluster, whereas the other half were assigned to the MD-dominant cluster. Out of the 91 possible pairs of the remaining 14 language fROIs, 86 pairs (94.5%) in the resting state condition and 91 pairs (100%) in the story-comprehension condition had their voxels jointly clustered significantly more often than expected by chance. Similarly, out of the 153 possible pairs of the 18 MD fROIs, 140 pairs (91.5%) in the resting state condition and 153 pairs (100%) in the story comprehension condition had their voxels jointly clustered significantly more often than expected by chance. Conversely, out of the 288 possible pairs consisting of a language fROI and a MD fROI, 286 pairs (99.3%) in the resting state condition and 287 pairs (99.7%) in the story-comprehension condition had their voxels jointly clustered less often than expected by chance. As was the case for fROI-wise clustering, the general patterns of voxel-wise clustering also did not depend on the choice of K (Figure 1.5).

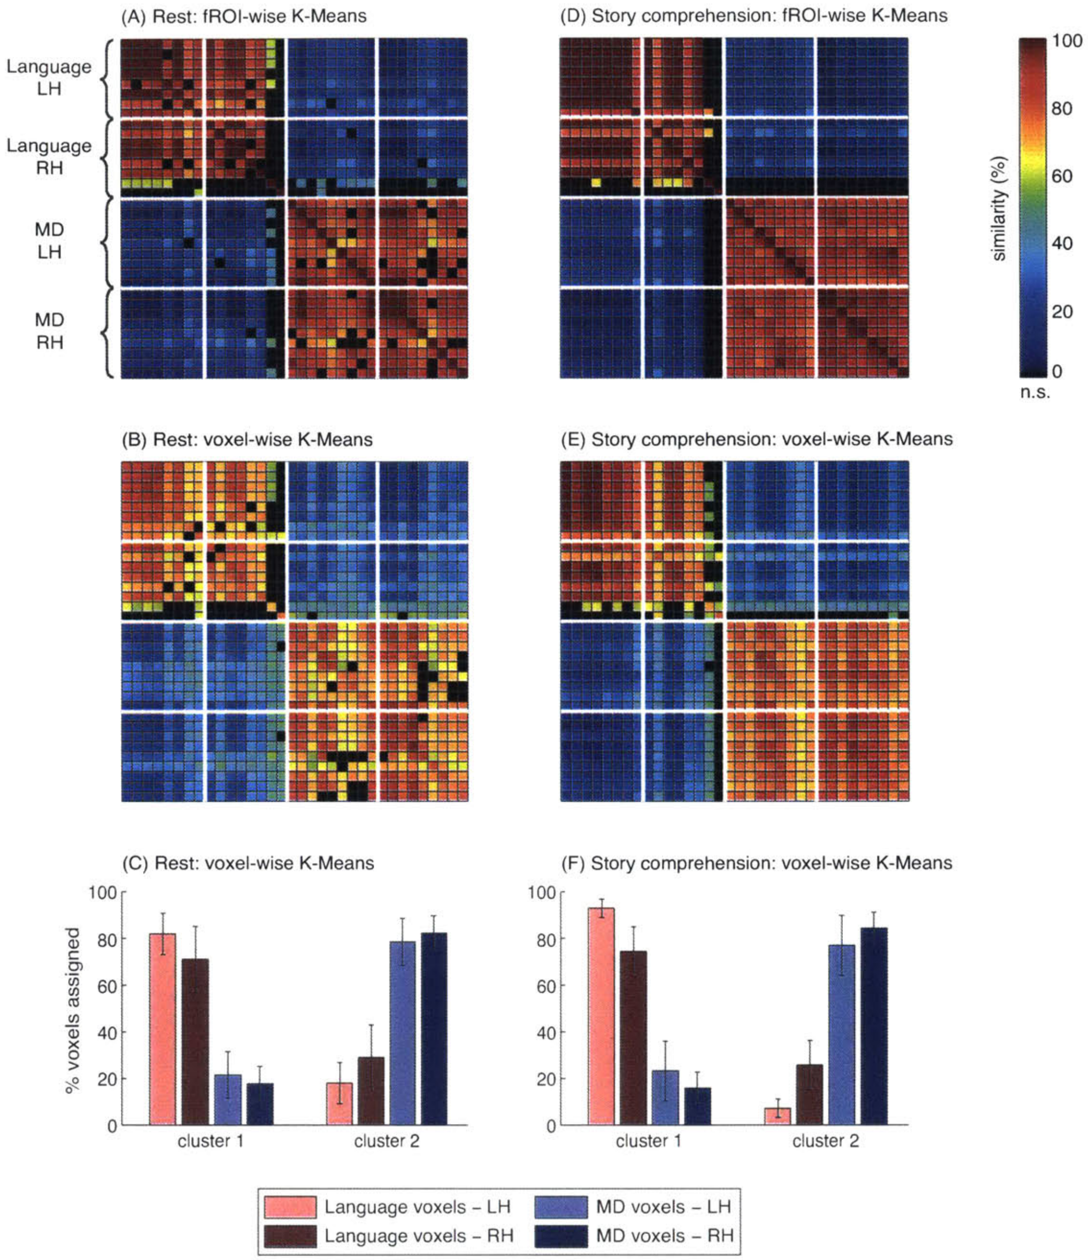

Figure 1.4. *K*-means clustering results for the resting state (left) and story comprehension (right) conditions, with $K=2$ clusters. (A), (D) The average BOLD signal time-course of each fROI was extracted, and the resulting time courses were clustered. In the fROI-to-ROI similarity matrices plotted here, the color of an entry (i,j) for a given pair of fROIs represents the probability (percentage of clustering solutions across participants and initializations) that the two fROIs would both be assigned to the same cluster. (B), (E) BOLD signal time-courses of all voxels falling within the fROIs were clustered. For each fROI, its "dominant cluster" was then defined as the cluster to which most of the voxels originating within that fROI were assigned. In the fROI-to-ROI similarity matrices plotted here, the color of an entry (i,j) for a given pair of fROIs represents the percentage of voxels in fROI j that were assigned to the dominant cluster of fROI i (note that this is not symmetrical). Percentages are averaged across participants and initializations. In all matrices (A)-(D) only significant entries are shown (as assessed with a permutation test, based on phase-shuffling of the original BOLD time-courses; $\alpha=0.05$, FDR-corrected). Non-significant entries are colored in black. The order of fROIs across rows (and columns) follows Table 1.1, where regions are sorted by system (language, then MD). Within each system, fROIs are sorted by hemisphere (left hemisphere (LH), then right hemisphere (RH)). Thick white lines separate these subsets of fROIs. (C), (F) Same data as in (B) and (E), respectively. The proportion of "language voxels" and "MD voxels" from each hemisphere that were assigned to each cluster are presented (across the two clusters, bars of the same color add to 100%). Errorbars show standard deviations across subjects.

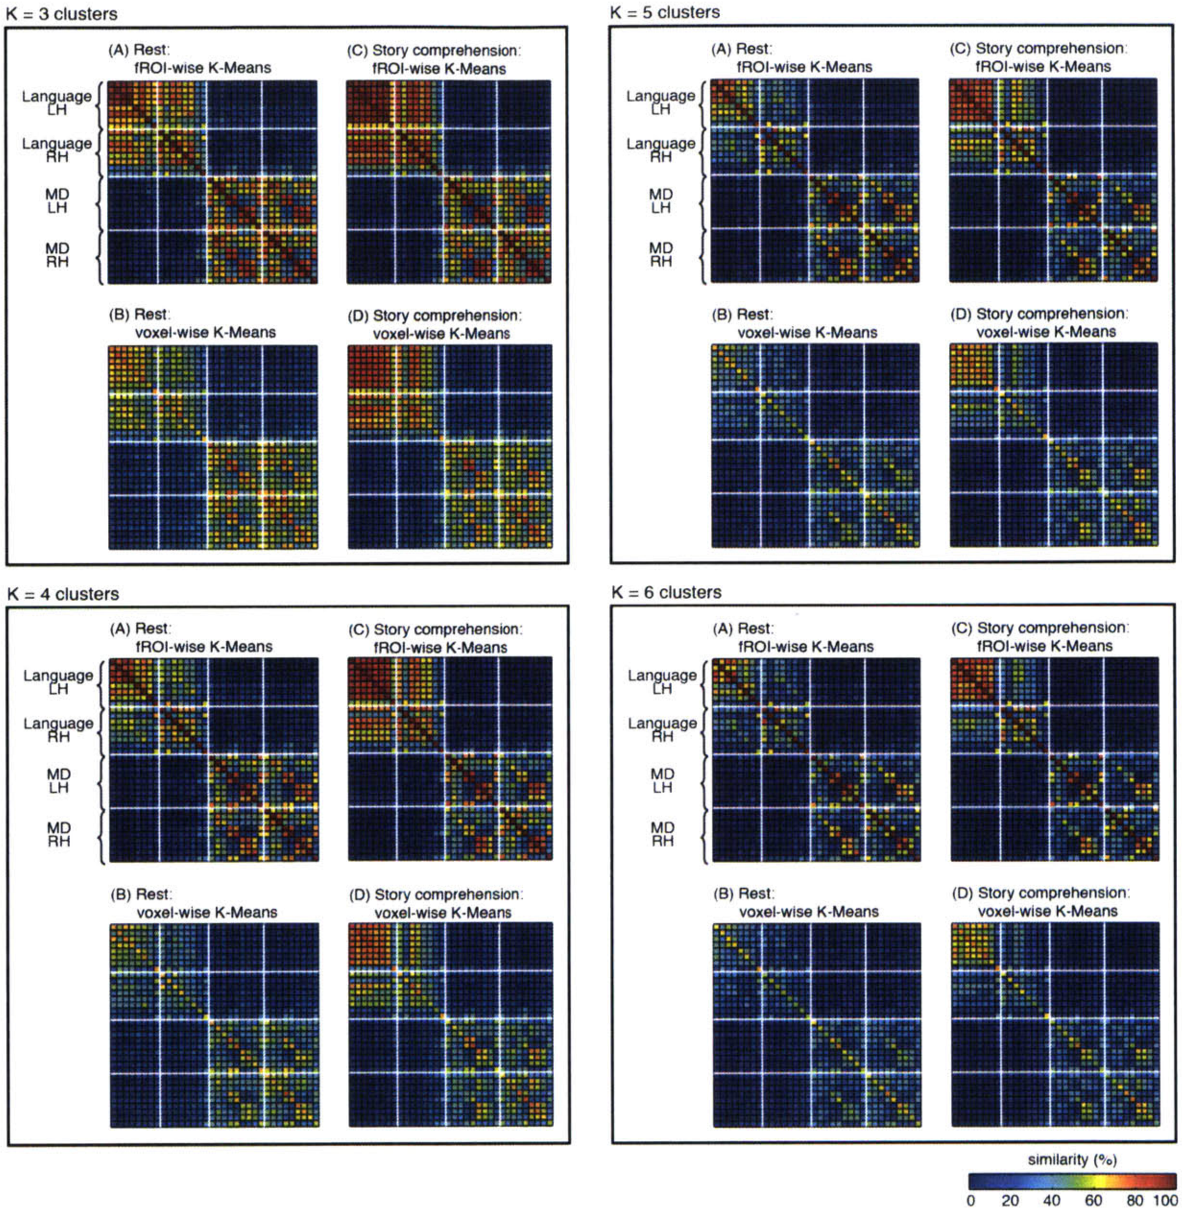

Figure 1.5. *K*-means clustering results of functional correlation data as a function of *K*. Conventions are the same as in Figure 1.4.

1.3.3. Hierarchical Clustering

Tree structures (dendrograms) of fROIs constructed based on their correlations revealed a functional architecture dominated by the partition into language and MD systems (Figure 1.6). In both the resting state and the story comprehension conditions, the topmost branching of the tree already separated the set of candidate language fROIs from the set of candidate MD fROIs. Thus, clustering together regions whose distance was shorter than the height of the first branching partitioned the data into language and MD systems. In both experimental conditions, this partition into 2 clusters had the highest modularity value compared to all other possible partitions licensed by the hierarchical tree, suggesting that the functional dissociation between the language and MD systems is the key organizational principle underlying the data.

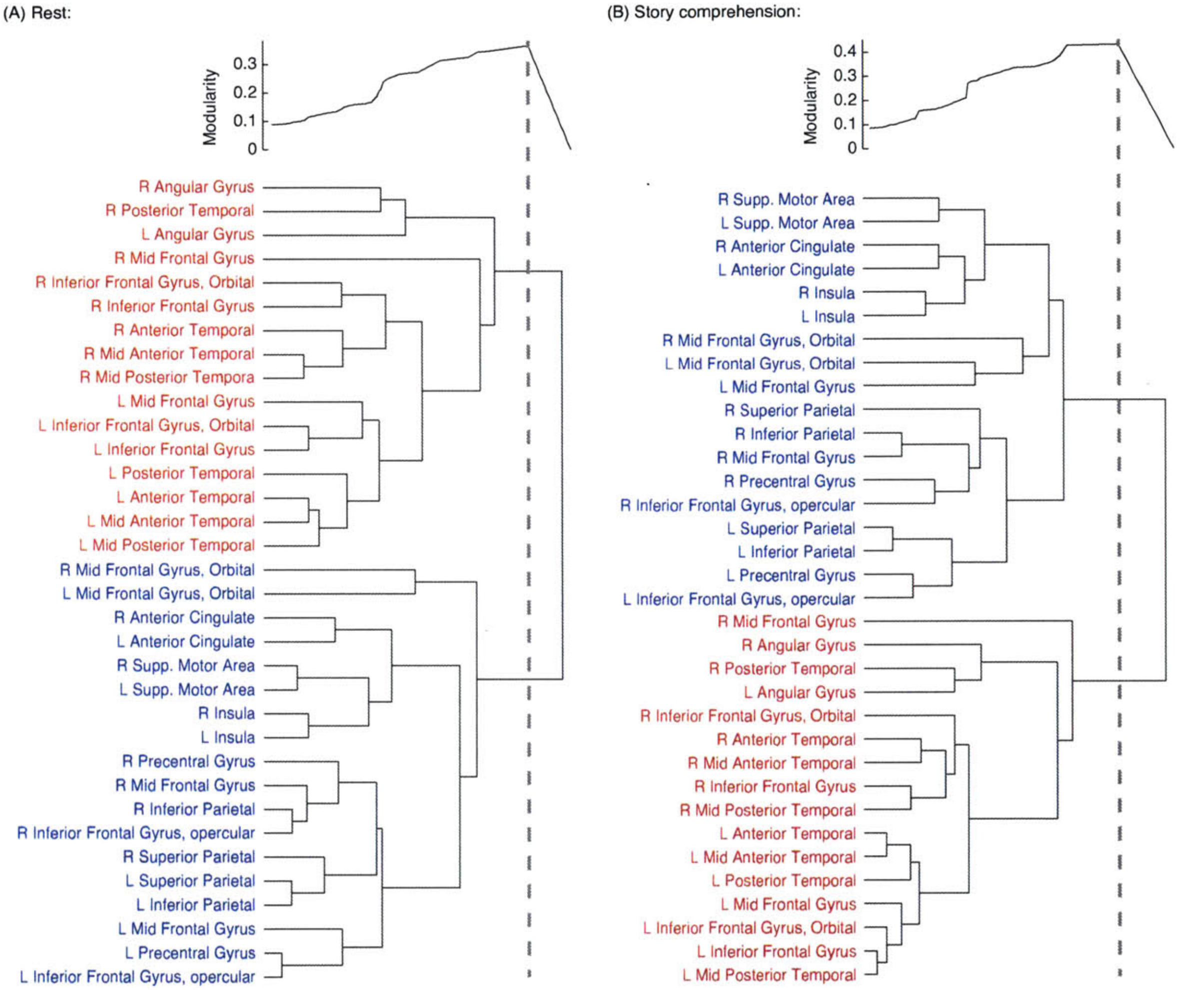

Figure 1.6. Hierarchical clustering results for (A) the resting state condition; and (B) the story comprehension condition. Hierarchical clustering creates a binary tree, with branch length (here, horizontal lines) corresponding to the similarity between fROIs (or sets of fROIs). Above each hierarchical tree, modularity is plotted for all fROI partitions licensed by the tree. Each point on the modularity plot corresponds to a partition generated by drawing an imaginary vertical line from that point through the tree, and clustering together only those fROIs that are merged to the left of this line (fROIs that are merged to the right of the line remain in separate clusters). A sample vertical line is drawn for the maximal modularity, which corresponds to a partition of the data into two clusters, one consisting of language fROIs and the other consisting of MD fROIs.

The organization of fROIs within each of the two systems provides further clues into their functional architecture. This organization could be characterized by using a short distance threshold for clustering regions into small subgroups, and then slowly increasing the threshold so that they merge into larger subgroups. The functional organization thus revealed was largely similar across the resting state and story comprehension conditions. In the language system, especially during rest, clustering was dominated by hemisphere and within each hemisphere regions were clustered according to their lobe. Namely, left temporal regions were clustered together, and only then merged with left frontal regions. Next, these left-hemisphere regions were merged with right-hemisphere regions, which themselves also formed temporal and frontal subsets. The right AngG and MFG were among the last to merge with the rest of the language system (consistent with our K-means results), along with the left AngG and right PostTemp.

In the MD system, clustering was sometimes dominated by hemisphere and sometimes by inter-hemispheric homology. Namely, most frontal and parietal regions tended to merge with each other ipsilaterally before merging across hemispheres. However, the Insula, SMA, ACC and MFG were each first clustered with their contralateral homologue, and only then merged with each other and with the former, fronto-parietal subset.

1.4. Discussion

The findings reported here demonstrate that fMRI BOLD signal fluctuations are strongly correlated among different brain regions within the language system, as well as among different regions within the MD system, but correlations across these two systems are weak or negative. These results are robust, generalizing across two conditions: rest, where signal fluctuations are not driven by an external task, and story comprehension, where signal fluctuations are task-evoked; and they arise in similar form from both hypothesis-driven and data-driven analyses. These data provide powerful new evidence that each of these systems forms a cohesive, integrated network, yet the two systems are functionally dissociated from each other.

Previous studies that used data-driven clustering of voxels across the brain based on resting state functional data provided the groundwork for the present study, and even revealed systems that approximately resemble the language and MD networks investigated here (Yeo et al., 2011; Lee et al., 2012; Mantini et al., 2013; Tie et al., 2014). However, the only way to link the clusters that emerged in those studies to the wealth of knowledge about the functions of different brain regions is through "reverse inference" based on stereotaxic coordinates (Poldrack, 2006). Such anatomy-based inferences are particularly challenging for the language and MD networks, because the mapping between stereotaxic coordinates and functional regions is degraded by both the high variability across participants in the anatomical locations of each region (Paus et al., 1996; Amunts et al., 1999; Tomaiuolo et al., 1999; Juch et al., 2005; Frost and

Goebel, 2012; Tahmasebi et al., 2012) and the frequent proximity of language and MD regions (e.g., Fedorenko et al., 2012a). Here we circumvented these problems by performing our clustering analyses on regions (and voxels) that were defined functionally within each participant, thereby allowing a direct interpretation of the resulting clusters in terms of specific functional hypotheses.

The combination of a participant-specific functional localization approach with an analysis of functional correlations has been previously applied to the ventral visual pathway (Turk-Browne et al., 2010; Zhu et al., 2011; Zhen et al., 2013) and other regions (Heinzle et al., 2012; Harmelech et al., 2013). In fact, this method was used by the first paper to report resting state functional correlations (Biswal et al., 1995). However, no prior study has used this method to study the functional relationship between the language and MD networks. Specifically, while a few prior functional correlation studies did use functional localizers for defining either language (Makuuchi and Friederici, 2013; Newman et al., 2013) or MD (Dosenbach et al., 2007; Seeley et al., 2007) regions, most of these studies used group-level analyses of the localizer data (or coordinates from prior studies; Turken and Dronkers, 2011), again with the potential shortcomings described above. The few studies that did define candidate language regions in individual participants (Hampson et al., 2002; Morgan et al., 2009) have focused on small subsets of the language network rather than providing wide coverage of regions of interest.

The current study therefore provides new support for the hypothesis that the language and MD networks are dissociable from each other and are recruited for distinct cognitive processes. Our findings complement prior evidence from neuroimaging studies that used standard functional contrasts (Monti et al., 2009; Fedorenko et al., 2011; Fedorenko et al., 2012a; Monti and Osherson, 2012; Monti et al., 2012), as well as evidence from double dissociations in the patient literature (Broca, 1861/2006; Wernicke, 1874/1969; Luria et al., 1965; Happé et al., 1999; Butterworth, 2000; Varley and Siegal, 2000; Peretz and Coltheart, 2003; Varley et al., 2005; Apperly et al., 2006; Klessinger et al., 2007; Bek et al., 2010).

1.4.1. Remaining questions

A crucial question for further investigation concerns the fine-grained functional organization within the language network and within the MD network. Although each network is highly integrated as indexed by the strong correlations among its constituent regions found here, further functional subdivisions within each network are likely. Indeed, our clustering results already capture some possible subdivisions within each network, and some of these appear to correspond to those suggested in prior studies. Namely, within the MD network, our hierarchical clustering analyses revealed two subsets that may correspond to the previously identified "fronto-parietal" and "cingulo-opercular" networks (Koechlin et al., 1999; Dosenbach et al., 2006; Dosenbach et al., 2007; Nomura et al., 2010; Power et al., 2011; Mantini et al., 2013). These two networks are hypothesized to be differently recruited for controlling task-relevant cognitive strategies (for a review, see Power and Petersen, 2013). Within the language network, our hierarchical clustering revealed frontal and temporal subsets in each hemisphere, consistent with previous results from both aphasic patients (e.g., Geschwind, 1970; Gorno-Tempini et al., 2004) and fMRI studies (e.g., Hagoort, 2003; Hagoort, 2005; Snijders et al., 2009; Tie et al., 2014). The current approach of combining functional correlation measures with functional localizers that target specific cognitive functions is likely to prove powerful in further elucidating these hypothesized subdivisions within the language and MD networks.

A second unanswered question concerns the neurobiological significance of functional correlations across brain regions. It has been suggested that these correlations may in part reflect (1) anatomical connectivity (direct and indirect); and/or (2) history of co-activation (for reviews and additional accounts, see He et al., 2008; Shmuel and Leopold, 2008; Deco et al., 2010; Schölvinck et al., 2010; Deco and Corbetta, 2011; Keller et al., 2011; Matsui et al., 2011; Deco et al., 2013). The extent to which the patterns of correlations reported here correspond to direct anatomical connections therefore remains to be discovered. Although evidence from diffusion imaging is generally consistent with resting state functional correlation measures (Hermundstad et al., 2013), neither is a perfect measure of structural connectivity (Uğurbil et al., 2013), posing substantial challenges for a definitive answer to this question. Moreover, some functional correlations are critically task-dependent (Hermundstad et al., 2013). Thus, although the functional dissociation between the language and MD networks generalized across resting state and story comprehension, it is possible that this dissociation would be modulated under other cognitive states.

Third, although our data indicate that language and MD regions are independently recruited during cognitive processing, this conclusion need not imply that the two networks can never be engaged simultaneously. Indeed, many previous fMRI studies have reported activations in MD regions during some language processing tasks, especially when such processing is effortful (Kuperberg et al., 2003; Rodd et al., 2005; Novais-Santos et al., 2007; January et al., 2009; Meltzer et al., 2010; McMillan et al., 2012; Nieuwland et al., 2012; Wild et al., 2012; McMillan et al., 2013). These findings suggest that the domain-general cognitive-control mechanisms associated with the MD network may play a role in language processing (Fedorenko, 2014; Fedorenko and Thompson-Schill, 2014) and, hence, that the MD and language systems may co-activate in some circumstances. Interactions between these two systems, however, may be more pronounced on a fast millisecond-level timescale, and therefore may not be detectable in the BOLD signal fluctuations measured here, given the low temporal resolution of this signal. Thus an important question for future research concerns the frequency, nature, and functional importance of interactions and co-activations of the language and MD networks.

1.4.2. Conclusion

Our results support a functional dissociation between the language and MD networks: each network is strongly correlated within itself, but pairs of regions straddling the two networks show weak (or negative) correlations. The robustness of this dissociation across conditions and analyses suggests that it reflects a deep principle of the functional organization of the human brain. Thus, the current data help resolve the controversy in the prior neuroimaging literature (Thompson-Schill et al., 2005; Blumstein and Amso, 2013) in favor of the hypothesis that at least some of the neural mechanisms used for high-level language processing are distinct from those that support other cognitive functions.

1.5. Acknowledgements

We thank Susan Whitfield-Gabrieli for valuable advice on connectivity analysis in the CONN toolbox, and Terri Lynn Scott for generating functional regions of interest in native functional space. The authors are also grateful to Anastasia Vishnevetsky, Steve Piantadosi and Ted Gibson for help in constructing and recording the materials for the story comprehension task, and to Eyal Dechter and Alex Kell for help with the script for the stories experiment. For technical support during scanning, the authors thank Christina Triantafyllou, Steve Shannon, and Sheeba Arnold. We also acknowledge the Athinoula A. Martinos Imaging Center at McGovern Institute for Brain Research, MIT. This research was supported by the following funding sources: grant EY13455 to NK, a grant from the Ellison Medical Foundation to NK, a seed grant from the Simons Center for the Social Brain to NK, and Eunice Kennedy Shriver NICHD Award K99 HD-057522 to EF.

References

Amunts K, Schleicher A, Bürgel U, Mohlberg H, Uylings H, Zilles K (1999) Broca's region revisited: Cytoarchitecture and intersubject variability. Journal of Comparative Neurology 412:319-341.

Apperly IA, Samson D, Carroll N, Hussain S, Humphreys G (2006) Intact first-and second-order false belief reasoning in a patient with severely impaired grammar. Social Neuroscience 1:334-348.

Bassett DS, Porter MA, Wymbs NF, Grafton ST, Carlson JM, Mucha PJ (2013) Robust detection of dynamic community structure in networks. Chaos: An Interdisciplinary Journal of Nonlinear Science 23:013142.

Behzadi Y, Restom K, Liau J, Liu TT (2007) A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. NeuroImage 37:90-101.

Bek J, Blades M, Siegal M, Varley RA (2010) Language and spatial reorientation: Evidence from severe aphasia. Journal of Experimental Psychology: Learning, Memory, and Cognition 36:646-658.

Benjamini Y, Yekutieli D (2001) The control of the false discovery rate in multiple testing under dependency. Ann Stat 29:1165-1188.

Biswal B, Zerrin Yetkin F, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo - planar MRI. Magnetic Resonance in Medicine 34:537-541.

Blumstein SE (2009) Reflections on the cognitive neuroscience of language. In: The Cognitive Neurosciences (Gazzaniga M, ed), pp 1235-1240. Cambridge, MA: MIT Press.

Blumstein SE, Amso D (2013) Dynamic functional organization of language: Insights from functional neuroimaging. Perspectives on Psychological Science 8:44-48.

Braver TS, Reynolds JR, Donaldson DI (2003) Neural mechanisms of transient and sustained cognitive control during task switching. Neuron 39:713-726.

Broca P (1861/2006) Comments regarding the seat of the faculty of spoken language, followed by an observation of aphemia (loss of speech). In: Broca's Region (Grodzinsky Y, Amunts K, eds), pp 291-304. New York, NY: Oxford University Press.

Butterworth B (2000) The Mathematical Brain. London, UK: Macmillan.

Chiarello C, Banich MT, Mack M (2003) Parallel systems for processing language: Hemispheric complementarity in the normal brain. In: Mind, Brain, and Language: Multidisciplinary Perspectives (Banich MT, Mack M, eds), pp 229-247. Mahwah, NJ: Lawrence Erlbaum.

Cole MW, Schneider W (2007) The cognitive control network: Integrated cortical regions with dissociable functions. NeuroImage 37:343-360.

Cordes D, Haughton VM, Arfanakis K, Wendt GJ, Turski PA, Moritz CH, Quigley MA, Meyerand ME (2000) Mapping functionally related regions of brain with functional connectivity MR imaging. American Journal of Neuroradiology 21:1636-1644.

Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, Quigley MA, Meyerand ME (2001) Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. American Journal of Neuroradiology 22:1326-1333.

Damasio AR (1992) Aphasia. New England Journal of Medicine 326:531-539.

Deco G, Corbetta M (2011) The dynamical balance of the brain at rest. The Neuroscientist 17:107-123.

Deco G, Jirsa VK, McIntosh AR (2010) Emerging concepts for the dynamical organization of resting-state activity in the brain. Nature Reviews Neuroscience 12:43-56.

Deco G, Jirsa VK, McIntosh AR (2013) Resting brains never rest: Computational insights into potential cognitive architectures. Trends in Neurosciences:268-274.

Dehaene S, Spelke E, Pinel P, Stanescu R, Tsivkin S (1999) Sources of mathematical thinking: Behavioral and brain-imaging evidence. Science 284:970-974.

Dosenbach NUF, Fair DA, Cohen AL, Schlaggar BL, Petersen SE (2008) A dual-networks architecture of top-down control. Trends in Cognitive Sciences 12:99-105.

Dosenbach NUF, Visscher KM, Palmer ED, Miezin FM, Wenger KK, Kang HC, Burgund ED, Grimes AL, Schlaggar BL, Petersen SE (2006) A core system for the implementation of task sets. Neuron 50:799-812.

Dosenbach NUF, Fair DA, Miezin FM, Cohen AL, Wenger KK, Dosenbach RAT, Fox MD, Snyder AZ, Vincent JL, Raichle ME (2007) Distinct brain networks for adaptive and stable task control in humans. Proceedings of the National Academy of Sciences 104:11073-11078.

Duncan J (2010) The multiple-demand (MD) system of the primate brain: Mental programs for intelligent behaviour. Trends in Cognitive Sciences 14:172-179.

Duncan J, Owen AM (2000) Common regions of the human frontal lobe recruited by diverse cognitive demands. Trends in Neurosciences 23:475-483.

Fedorenko E (2014) The role of domain-general cognitive control in language comprehension. Frontiers in Psychology 5:335.

Fedorenko E, Thompson-Schill SL (2014) Reworking the language network. Trends in Cognitive Sciences 18:120-126.

Fedorenko E, Behr MK, Kanwisher N (2011) Functional specificity for high-level linguistic processing in the human brain. Proceedings of the National Academy of Sciences 108:16428-16433.

Fedorenko E, Duncan J, Kanwisher N (2012a) Language-selective and domain-general regions lie side by side within Broca's area. Current Biology.

Fedorenko E, Nieto-Castañón A, Kanwisher N (2012b) Lexical and syntactic representations in the brain: an fMRI investigation with multi-voxel pattern analyses. Neuropsychologia 50:499-513.

Fedorenko E, Duncan J, Kanwisher N (2013) Broad domain generality in focal regions of frontal and parietal cortex. Proceedings of the National Academy of Sciences 110:16616-16621.

Fedorenko E, Hsieh P-J, Nieto-Castañón A, Whitfield-Gabrieli S, Kanwisher N (2010) New method for fMRI investigations of language: Defining ROIs functionally in individual subjects. Journal of Neurophysiology 104:1177-1194.

Frost MA, Goebel R (2012) Measuring structural–functional correspondence: Spatial variability of specialised brain regions after macro-anatomical alignment. NeuroImage 59:1369-1381.

Geschwind N (1970) The organization of language and the brain. Science 170:940-944.

Gómez S, Jensen P, Arenas A (2009) Analysis of community structure in networks of correlated data. Physical Review E 80:016114.

Gorno - Tempini ML, Dronkers NF, Rankin KP, Ogar JM, Phengrasamy L, Rosen HJ, Johnson JK, Weiner MW, Miller BL (2004) Cognition and anatomy in three variants of primary progressive aphasia. Ann Neurol 55:335-346.

Hagoort P (2003) How the brain solves the binding problem for language: A neurocomputational model of syntactic processing. NeuroImage 20:S18-S29.

Hagoort P (2005) On Broca, brain, and binding: A new framework. Trends in Cognitive Sciences 9:416-423.

Hampshire A, Highfield RR, Parkin BL, Owen AM (2012) Fractionating human intelligence. Neuron 76:1225-1237.

Hampson M, Peterson BS, Skudlarski P, Gatenby JC, Gore JC (2002) Detection of functional connectivity using temporal correlations in MR images. Human Brain Mapping 15:247-262.

Happé F, Brownell H, Winner E (1999) Acquired "theory of mind" impairments following stroke. Cognition 70:211-240.

Harmelech T, Preminger S, Wertman E, Malach R (2013) The day-after effect: Long term, Hebbian-like restructuring of resting-state fMRI patterns induced by a single epoch of cortical activation. The Journal of Neuroscience 33:9488-9497.

Hartigan JA (1975) Clustering Algorithms. New York, NY: John Wiley & Sons.

He BJ, Snyder AZ, Zempel JM, Smyth MD, Raichle ME (2008) Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. Proceedings of the National Academy of Sciences 105:16039-16044.

Hein G, Knight RT (2008) Superior temporal sulcus—it's my area; or is it? Journal of Cognitive Neuroscience 20:2125-2136.

Heinzle J, Wenzel MA, Haynes J-D (2012) Visuomotor functional network topology predicts upcoming tasks. The Journal of Neuroscience 32:9960-9968.

Hermundstad AM, Bassett DS, Brown KS, Aminoff EM, Clewett D, Freeman S, Frithsen A, Johnson A, Tipper CM, Miller MB (2013) Structural foundations of resting-state and task-based functional connectivity in the human brain. Proceedings of the National Academy of Sciences 110:6169-6174.

January D, Trueswell JC, Thompson-Schill SL (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 21:2434-2444.

Jenkinson M, Bannister P, Brady M, Smith S (2002) Improved optimization for the robust and accurate linear registration and motion correction of brain images. NeuroImage 17:825-841.

Juch H, Zimine I, Seghier ML, Lazeyras F, Fasel JH (2005) Anatomical variability of the lateral frontal lobe surface: Implication for intersubject variability in language neuroimaging. NeuroImage 24:504-514.

Julian JB, Fedorenko E, Webster J, Kanwisher N (2012) An algorithmic method for functionally defining regions of interest in the ventral visual pathway. NeuroImage 60:2357-2364.

Jung-Beeman M (2005) Bilateral brain processes for comprehending natural language. Trends in Cognitive Sciences 9:512-518.

Kaan E, Swaab TY (2002) The brain circuitry of syntactic comprehension. Trends in Cognitive Sciences 6:350-356.

Keller CJ, Bickel S, Entz L, Ulbert I, Milham MP, Kelly C, Mehta AD (2011) Intrinsic functional architecture predicts electrically evoked responses in the human brain. Proceedings of the National Academy of Sciences 108:10308-10313.

Klessinger N, Szczerbinski M, Varley RA (2007) Algebra in a man with severe aphasia. Neuropsychologia 45:1642-1648.

Koechlin E, Jubault T (2006) Broca's area and the hierarchical organization of human behavior. Neuron 50:963-974.

Koechlin E, Basso G, Pietrini P, Panzer S, Grafman J (1999) The role of the anterior prefrontal cortex in human cognition. Nature 399:148-151.

Koelsch S, Gunter TC, Zysset S, Lohmann G, Friederici AD (2002) Bach speaks: A cortical “language-network” serves the processing of music. NeuroImage 17:956-966.

Kuperberg GR, Holcomb PJ, Sitnikova T, Greve D, Dale AM, Caplan D (2003) Distinct patterns of neural modulation during the processing of conceptual and syntactic anomalies. Journal of Cognitive Neuroscience 15:272-293.

Lancichinetti A, Fortunato S (2012) Consensus clustering in complex networks. Scientific Reports 2:336.

Lee MH, Hacker CD, Snyder AZ, Corbetta M, Zhang D, Leuthardt EC, Shimony JS (2012) Clustering of resting state networks. PloS one 7:e40370.

Lemieux L, Salek-Haddadi A, Lund TE, Laufs H, Carmichael D (2007) Modelling large motion events in fMRI studies of patients with epilepsy. Magnetic Resonance Imaging 25:894-901.

Luria AR, Tsvetkova LS, Futer DS (1965) Aphasia in a composer. Journal of the Neurological Sciences 2:288-292.

Maess B, Koelsch S, Gunter TC, Friederici AD (2001) Musical syntax is processed in Broca's area: an MEG study. Nature Neuroscience 4:540-545.

Makuuchi M, Friederici AD (2013) Hierarchical functional connectivity between the core language system and the working memory system. Cortex 49:2416-2423.

Mantini D, Corbetta M, Romani GL, Orban GA, Vanduffel W (2013) Evolutionarily novel functional networks in the human brain? The Journal of Neuroscience 33:3259-3275.

Matsui T, Tamura K, Koyano KW, Takeuchi D, Adachi Y, Osada T, Miyashita Y (2011) Direct comparison of spontaneous functional connectivity and effective connectivity measured by intracortical microstimulation: an fMRI study in macaque monkeys. Cerebral Cortex 21:2348-2356.

McMillan CT, Clark R, Gunawardena D, Ryant N, Grossman M (2012) fMRI evidence for strategic decision-making during resolution of pronoun reference. Neuropsychologia 50:674-687.

McMillan CT, Coleman D, Clark R, Liang T-W, Gross RG, Grossman M (2013) Converging evidence for the processing costs associated with ambiguous quantifier comprehension. Frontiers in Psychology 4:153.

Meltzer JA, McArdle JJ, Schafer RJ, Braun AR (2010) Neural aspects of sentence comprehension: syntactic complexity, reversibility, and reanalysis. Cerebral Cortex 20:1853-1864.

Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. Annual Review of Neuroscience 24:167-202.

Monti MM, Osherson DN (2012) Logic, language and the brain. Brain Research 1428:33-42.

Monti MM, Parsons LM, Osherson DN (2009) The boundaries of language and thought in deductive inference. Proceedings of the National Academy of Sciences 106:12554-12559.

Monti MM, Parsons LM, Osherson DN (2012) Thought Beyond Language Neural Dissociation of Algebra and Natural Language. Psychological Science 23:914-922.

Morgan VL, Mishra A, Newton AT, Gore JC, Ding Z (2009) Integrating functional and diffusion magnetic resonance imaging for analysis of structure-function relationship in the human language network. PloS one 4:e6660.

Newman AJ, Kenny S, Saint-Aubin J, Klein RM (2013) Can skilled readers perform a second task in parallel? A functional connectivity MRI study. Brain and Language 124:84-95.

Newman MEJ, Girvan M (2004) Finding and evaluating community structure in networks. Physical Review E 69:026113.

Nieuwland MS, Martin AE, Carreiras M (2012) Brain regions that process case: Evidence from Basque. Human Brain Mapping 33:2509-2520.

Nomura EM, Gratton C, Visser RM, Kayser A, Perez F, D'Esposito M (2010) Double dissociation of two cognitive control networks in patients with focal brain lesions. Proceedings of the National Academy of Sciences 107:12017-12022.

Novais-Santos S, Gee J, Shah M, Troiani V, Work M, Grossman M (2007) Resolving sentence ambiguity with planning and working memory resources: Evidence from fMRI. NeuroImage 37:361-378.

Paus T, Tomaiuolo F, Otaky N, MacDonald D, Petrides M, Atlas J, Morris R, Evans AC (1996) Human cingulate and paracingulate sulci: Pattern, variability, asymmetry, and probabilistic map. Cerebral Cortex 6:207-214.

Peretz I, Coltheart M (2003) Modularity of music processing. Nature Neuroscience 6:688-691.

Poldrack RA (2006) Can cognitive processes be inferred from neuroimaging data? Trends in Cognitive Sciences 10:59-63.

Postelnicu G, Zollei L, Fischl B (2009) Combined volumetric and surface registration. IEEE Transactions on Medical Imaging 28:508-522.

Power JD, Petersen SE (2013) Control-related systems in the human brain. Current Opinion in Neurobiology 23:223-228.

Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE (2012) Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. NeuroImage 59:2142-2154.

Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL (2011) Functional network organization of the human brain. Neuron 72:665-678.

Rodd JM, Davis MH, Johnsrude IS (2005) The neural mechanisms of speech comprehension: fMRI studies of semantic ambiguity. Cerebral Cortex 15:1261-1269.

Satterthwaite TD, Wolf DH, Loughead J, Ruparel K, Elliott MA, Hakonarson H, Gur RC, Gur RE (2012) Impact of in-scanner head motion on multiple measures of functional connectivity: Relevance for studies of neurodevelopment in youth. NeuroImage 60:623-632.

Satterthwaite TD, Elliott MA, Gerraty RT, Ruparel K, Loughead J, Calkins ME, Eickhoff SB, Hakonarson H, Gur RC, Gur RE (2013) An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. NeuroImage 64:240-256.

Schölvinck ML, Maier A, Frank QY, Duyn JH, Leopold DA (2010) Neural basis of global resting-state fMRI activity. Proceedings of the National Academy of Sciences 107:10238-10243.

Seeley WW, Menon V, Schatzberg AF, Keller J, Glover GH, Kenna H, Reiss AL, Greicius MD (2007) Dissociable intrinsic connectivity networks for salience processing and executive control. The Journal of Neuroscience 27:2349-2356.

Shmuel A, Leopold DA (2008) Neuronal correlates of spontaneous fluctuations in fMRI signals in monkey visual cortex: implications for functional connectivity at rest. Human Brain Mapping 29:751-761.

Silver NC, Dunlap WP (1987) Averaging correlation coefficients: Should Fisher's z transformation be used? Journal of Applied Psychology 72:146.

Snijders TM, Vosse T, Kempen G, Van Berkum JJA, Petersson KM, Hagoort P (2009) Retrieval and unification of syntactic structure in sentence comprehension: an fMRI study using word-category ambiguity. Cerebral Cortex 19:1493-1503.

Stanescu-Cosson R, Pinel P, van de Moortele P-F, Le Bihan D, Cohen L, Dehaene S (2000) Understanding dissociations in dyscalculia: A brain imaging study of the impact of number size on the cerebral networks for exact and approximate calculation. Brain 123:2240-2255.

Tahmasebi AM, Davis MH, Wild CJ, Rodd JM, Hakyemez H, Abolmaesumi P, Johnsrude IS (2012) Is the link between anatomical structure and function equally strong at all cognitive levels of processing? Cerebral Cortex 22:1593-1603.

Thesen S, Heid O, Mueller E, Schad LR (2000) Prospective acquisition correction for head motion with image - based tracking for real - time fMRI. Magnetic Resonance in Medicine 44:457-465.

Thompson-Schill SL, Bedny M, Goldberg RF (2005) The frontal lobes and the regulation of mental activity. Current Opinion in Neurobiology 15:219-224.

Tie Y, Rigolo L, Norton IH, Huang RY, Wu W, Orringer D, Mukundan S, Golby AJ (2014) Defining language networks from resting - state fMRI for surgical planning–A feasibility study. Human Brain Mapping 35:1018-1030.

Tomaiuolo F, MacDonald JD, Caramanos Z, Posner G, Chiavaras M, Evans AC, Petrides M (1999) Morphology, morphometry and probability mapping of the pars opercularis of the inferior frontal gyrus: An in vivo MRI analysis. European Journal of Neuroscience 11:3033-3046.

Turk-Browne NB, Norman-Haignere SV, McCarthy G (2010) Face-specific resting functional connectivity between the fusiform gyrus and posterior superior temporal sulcus. Frontiers in Human Neuroscience 4:176.

Turken U, Dronkers NF (2011) The neural architecture of the language comprehension network: Converging evidence from lesion and connectivity analyses. Frontiers in Systems Neuroscience 5:1.

Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M (2002) Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. NeuroImage 15:273-289.

Uğurbil K, Xu J, Auerbach EJ, Moeller S, Vu A, Duarte-Carvajalino JM, Lenglet C, Wu X, Schmitter S, Van de Moortele PF (2013) Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. NeuroImage 80:80-104.

Van Dijk KRA, Hedden T, Venkataraman A, Evans KC, Lazar SW, Buckner RL (2010) Intrinsic functional connectivity as a tool for human connectomics: Theory, properties, and optimization. Journal of Neurophysiology 103:297-321.

Varley RA, Siegal M (2000) Evidence for cognition without grammar from causal reasoning and "theory of mind" in an agrammatic aphasic patient. Current Biology 10:723-726.

Varley RA, Klessinger NJC, Romanowski CAJ, Siegal M (2005) Agrammatic but numerate. Proceedings of the National Academy of Sciences 102:3519-3524.

Wernicke C (1874/1969) The symptom complex of aphasia: A psychological study on an anatomical basis. In: Boston Studies in the Philosophy of Science (Cohen RS, Wartofsky MW, eds), pp 34-97. Dordrecht, The Netherlands: D. Reidel Publishing.

Whitfield-Gabrieli S, Nieto-Castañón A (2012) Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. Brain Connectivity 2:125-141.

Wild CJ, Yusuf A, Wilson DE, Peelle JE, Davis MH, Johnsrude IS (2012) Effortful listening: the processing of degraded speech depends critically on attention. The Journal of Neuroscience 32:14010-14021.

Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of Neurophysiology 106:1125-1165.

Yue Q, Zhang L, Xu G, Shu H, Li P (2013) Task-modulated activation and functional connectivity of the temporal and frontal areas during speech comprehension. Neuroscience.

Zhen Z, Fang H, Liu J (2013) The hierarchical brain network for face recognition. PloS One 8:e59886.

Zhu Q, Zhang J, Luo YLL, Dilks DD, Liu J (2011) Resting-state neural activity across face-selective cortical regions is behaviorally relevant. The Journal of Neuroscience 31:10323-10330.

[This page intentionally left blank]

Chapter 2

Close tracking of linguistic input by language-selective but not domain-general brain regions

With Evelina Fedorenko

Abstract

Language comprehension engages a cortical network of left frontal and temporal regions (Bates et al., 2003; Jung-Beeman, 2005; Fedorenko et al., 2010; Menenti et al., 2011; Mesulam et al., 2015; Mirman et al., 2015). Activity in this network is sensitive to linguistic features such as lexical information, syntax and compositional semantics (Keller et al., 2001; Fedorenko et al., 2012b; Bautista and Wilson, 2016; Blank et al., 2016). However, this network shows virtually no engagement in non-linguistic tasks (Monti et al., 2009; Fedorenko et al., 2011; Monti and Osherson, 2012; Monti et al., 2012) and is therefore language-selective. In addition, language comprehension engages a second network consisting of frontal, parietal, cingulate, and insular regions (Thompson-Schill et al., 2005; Bedny et al., 2008; Hindy et al., 2012; Hindy et al., 2015). Activity in this "Multiple Demand (MD)" network (Duncan, 2010) is sensitive to comprehension difficulty, increasing in the presence of e.g. ambiguity (Rodd et al., 2005; Nieuwland et al., 2007; Novais-Santos et al., 2007; January et al., 2009; Ye and Zhou, 2009; McMillan et al., 2012; McMillan et al., 2013), infrequent words (Fiez et al., 1999; Fiebach et al., 2002; Chee et al., 2003; Nakic et al., 2006; Hauk et al., 2008; Yarkoni et al., 2008; Carreiras et al., 2009) and non-local syntactic dependencies (Stromswold et al., 1996; Stowe et al., 1998; Caplan et al., 1999; Constable et al., 2004; Chen et al., 2006; Peelle et al., 2009; Barde et al., 2012). However, this network similarly scales its activity with cognitive effort across a wide range of non-linguistic tasks (Duncan, 2010; Fedorenko et al., 2013) and is therefore domain-general. Given the functional dissociation between the language and MD networks (Fedorenko et al., 2012a; Blank et al., 2014), their respective contributions to comprehension are likely distinct, yet such

differences remain elusive. Critically, given that each network is sensitive to some linguistic features, prior research has presupposed that both networks track linguistic input closely, and in a manner consistent across individuals. Here, we used fMRI to test this assumption by comparing the BOLD signal time-courses in each network across different individuals listening to the same story (Wilson et al., 2008; Lerner et al., 2011; Honey et al., 2012). Language network activity showed fewer individual differences, indicative of closer input tracking, whereas MD network activity was more idiosyncratic and, moreover, showed low reliability within an individual across repetitions of a story. These findings constrain cognitive models of language comprehension by suggesting a novel distinction between the processes implemented in the language and MD networks.

2.1. Results

2.1.1. Correlations in network activity across individuals listening to a story

To characterize the functional contributions of the language and MD networks to comprehension, we tested how closely each network tracks naturalistic linguistic input (stories) and how similar such tracking is across individuals. Our interest in naturalistic cognition is three-fold: first, some brain regions respond more reliably to richly structured natural input compared to experimentally controlled input (Hasson et al., 2010). Second, task-free natural language processing plausibly differs from task-based processing, especially given that MD regions operate in a task-dependent manner by biasing representations in other cortical networks in favor of task-relevant features (Miller and Cohen, 2001; Sreenivasan et al., 2014; D'Esposito and Postle, 2015). Third, naturalistic stories require all aspects of the linguistic input to be combined into a single, rich representation, unlike experimental stimuli that focus on particular linguistic features and have lower ecological validity.

Prior to the story comprehension scan, language and MD regions were functionally identified in each individual participant. Language regions were localized using a reading task that contrasted sentences with pronounceable nonwords (Figure 2.1A). We localized 8 left-hemispheric (LH) regions (Figure 2.2A) as well as 8 right-hemispheric (RH) homologues, which are also activated during some aspects of language processing (George et al., 1996; Binder et al., 1997; Chiarello et al., 2003; Jung-Beeman, 2005; Fedorenko et al., 2010; Price, 2012; Yang, 2014; Bozic et al., 2015) but might differ from LH regions in their contribution to comprehension, as suggested by neuroimaging (Fedorenko and Thompson-Schill, 2014) and neuropsychological data (Geschwind, 1970; Damasio, 1992). MD regions were functionally identified using a spatial working memory task that contrasted a hard version with an easy version (Figure 2.1B). We localized nine regions in each hemisphere (Figure 2.2B) and, based on prior findings (Dosenbach et al., 2006; Dosenbach et al., 2007; Nomura et al., 2010; Power et al., 2011; Mantini et al., 2013), grouped them into two functionally distinct sub-networks: fronto-parietal (MDfp) and cingulo-opercular (MDco) (similar results were obtained when regions were instead grouped by hemisphere).

Each participant (*n*=19) then listened to 1-4 stories (duration: 270s-364s) constructed from publicly available texts, each followed by a comprehension test. **To ensure that the stories strongly engaged the MD network, they were edited to include frequent occurrences of linguistic phenomena that increase processing difficulty and have been demonstrated to recruit this network (Figure 2.1C) (such phenomena are not naturally frequent enough; Collins, 1996; Roland et al., 2007; Futrell et al., 2015). Following (Honey et al., 2012),** we reasoned that if a given brain region closely tracked linguistic input with little individual differences then its activity time-course would be similar across participants and would thus show high inter-subject correlations (ISCs) (Hasson et al., 2004). Therefore, we recorded the BOLD signal time-course in each language and MD region during each story and computed regional ISCs. To ensure that ISCs reflected tracking of linguistic information rather than low-level sensory information, signals were first regressed against time-courses from the auditory cortex (similar results were obtained without regression).

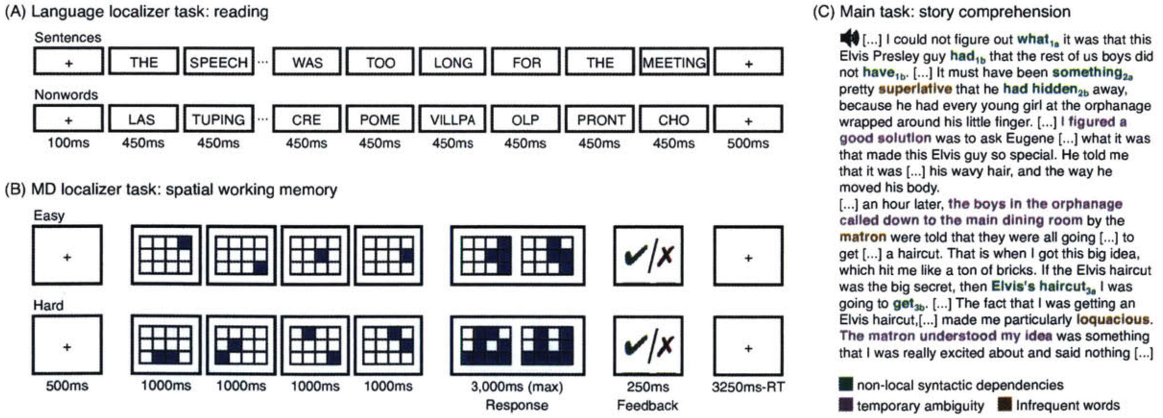

[...] I could not figure out **what$_{1a}$** it was that this Elvis Presley guy **had$_{1b}$** that the rest of us boys did not **have$_{1b}$**. [...] It must have been **something$_{2a}$** pretty **superlative** that he **had hidden$_{2b}$** away, because he had every young girl at the orphanage wrapped around his little finger. [...] **I figured a good solution** was to ask Eugene [...] what it was that made this Elvis guy so special. He told me that it was [...] his wavy hair, and the way he moved his body.
[...] an hour later, **the boys in the orphanage called down to the main dining room** by the **matron** were told that they were all going [...] to get [...] a haircut. That is when I got this big idea, which hit me like a ton of bricks. If the Elvis haircut was the big secret, then **Elvis's haircut$_{3a}$** I was going to **get$_{3b}$**. [...] The fact that I was getting an Elvis haircut,[...] made me particularly **loquacious**. **The matron understood my idea** was something that I was really excited about and said nothing [...]

■ non-local syntactic dependencies
■ temporary ambiguity ■ Infrequent words

Figure 2.1. Experimental tasks. (A) The reading task used to localize language regions, based on the critical contrast *sentences* > *nonwords*. (B) The spatial working memory task used to localize MD regions, based on the critical contrast *hard* > *easy*. (C) An excerpt from a story used in the main comprehension experiment. Linguistic phenomena that increase processing difficulty and have been shown to recruit the MD network, but are naturally infrequent, were edited into the text. These include non-local syntactic dependencies (green; words in this relation have subscripts with the same number but different letters); temporary ambiguity (purple), where a likely initial parse is later revealed to be wrong; and low frequency words (brown).

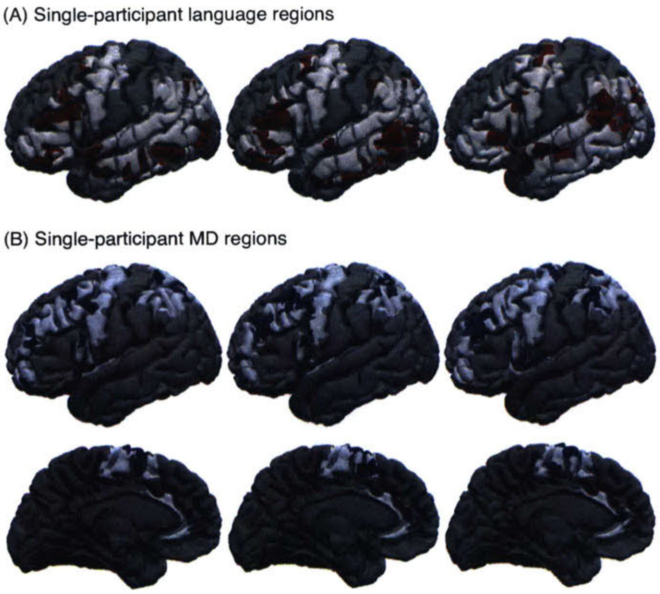

(A) Single-participant language regions

(B) Single-participant MD regions

Figure 2.2. Functional regions of the language and MD networks. (A) LH language regions in 3 individual participants are shown in dark red. These regions were localized with a reading task (see Figure 2.1A). They were constrained to fall within eight broad areas where activations for this task are common across the population, shown in light pink. These areas were defined based on group-level data from a previous sample (Fedorenko et al., 2010). (B) LH MD regions of the same 3 participants are shown in dark blue. These regions were localized with a spatial working memory task (see Figure 2.1B). They were constrained to fall within nine broad areas where activations for this localizer are common across the population, shown in light blue. These areas were anatomically defined (Fedorenko et al., 2013).

We used linear, mixed-effect models to test whether the language and MD networks differed from each other in the degree of stimulus tracking, as estimated via their ISCs. Across stories, the LH language network showed the highest ISCs (Fisher transformed r=0.280), stronger than ISCs in the RH language network (r=0.210, $p<10^{-9}$), the MDfp network (r=0.136, $p\approx0$) and the MDco network (r=0.117, $p\approx0$). The RH language network, in turn, showed higher ISCs than both the MDfp network ($p<10^{-11}$) and the MDco network ($p<10^{-13}$). The two MD networks did not differ from each other (p=0.218) (Figure 2.3; all p-values here and elsewhere are corrected for multiple comparisons using false-discovery rate (FDR) correction; Benjamini and Yekutieli, 2001). The difference between the LH language network and the two MD networks was also observed for each story separately.

Next, we performed two replication experiments. In the first (n=13), we computed ISCs in a new participant group for two of the stories used above. In the second (n=19), we used a story recorded at a live event (duration: 420s; Lerner et al., 2011) that did not undergo linguistic editing and was thus even more naturalistic than our previous stimuli. In both experiments, we again found that ISCs in the LH language network (replication 1: r=0.252; replication 2: r=0.303) were stronger than in the RH language network (r=0.172, $p<10^{-7}$; r=0.250, p=0.001), the MDfp network (r=0.147, $p<10^{-15}$; r=0.160, $p\approx0$) and the MDco network (r=0.114, $p\approx0$; r=0.163, $p<10^{-15}$). ISCs in the RH language network were somewhat stronger than ISCs in the MDfp network (p=0.066; $p<10^{-9}$) and stronger than ISCs in the MDco network ($p<0.001$; $p<10^{-7}$). The two latter networks reliably differed from each other only in the first replication ($p<0.033$) (Figure 2.3).

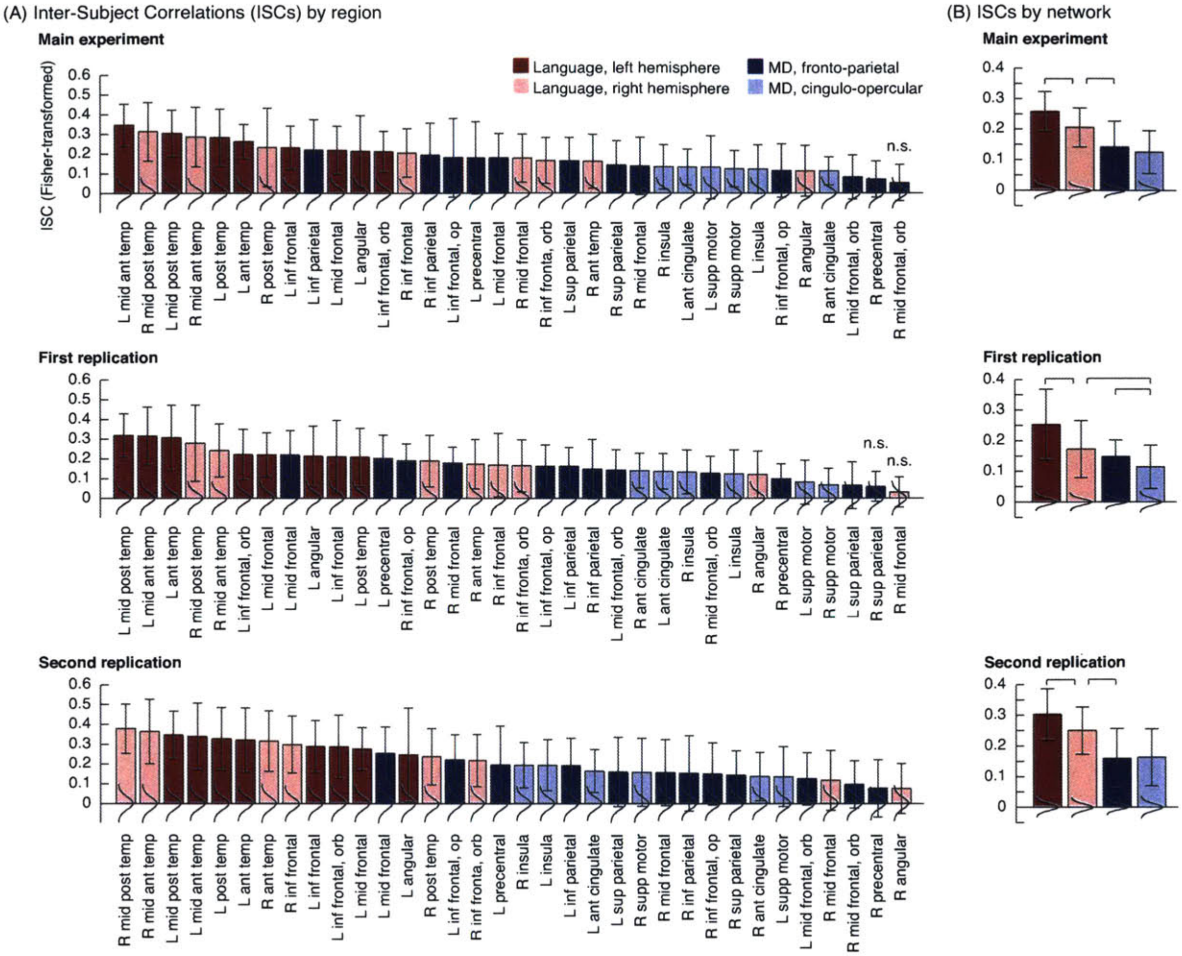

Figure 2.3. ISCs during story comprehension in the language and MD networks. (A) Bar plots of (Fisher-transformed) ISCs for each brain region. Error bars show standard deviations across participants. Vertical curves show Gaussian fits to empirical null distributions against which each bar can be tested (n.s., non-significant results at a threshold of 0.05; FDR corrected). Regions are ordered according to ISC magnitude in the main experiment (top), first replication (middle), and second replication (bottom), and are grouped into 4 functional networks indicated by bar color. Whereas the ordering of individual regions is not necessarily reliable across experiments, a replicable pattern emerges where ISCs are stronger in language regions (red) than in MD regions (blue). (B) Mean ISCs within each functional network, same conventions as in (A). Horizontal lines connect pairs of bars that significantly differ (in each pair, the left ISC is greater than the right ISCs and all ISCs that are further to the right).

Across these three experiments, we find that signals in the language and MD networks differ in their ISCs and, thus, in the percentage of variance they share across people. To further interpret these findings we computed an "upper bound" on ISCs, reflecting the highest values that could be expected in our measurements; namely, we computed ISCs in low-level auditory regions (defined anatomically) that track sensory input very closely (Lerner et al., 2011). Combining data across experiments, these auditory ISCs are estimated at r=0.450. Thus, signals in the LH language network (r=0.287) share 40.8% of this "maximum shareable variance" across people; signals in the RH language network (r=0.216) share 23%, whereas signals in the MDfp network

(r=0.153) and MDco network (r=0.134) share only 11.6% and 8.8%, respectively. Importantly, however, almost all ISCs – even those in MD regions – are significantly greater than expected by chance (Figure 2.3). Therefore, even domain-general MD regions track stories to a non-trivial extent.

2.1.2. Correlations in network activity within individuals listening to a story twice

The relatively low ISCs in MD regions could be interpreted in two ways: on the one hand, MD regions might closely track linguistic input but do so in an idiosyncratic fashion across individuals. For example, if different people find different sections of the story difficult to comprehend, they might each recruit their MD network at respectively different times. In this case, MD activity time-courses would be stimulus-related for each individual but would differ across individuals. Alternatively, activity in the MD regions might not be closely linked to the linguistic input at all. These two interpretations can be distinguished by correlating signal time-courses within a given individual who is listening to the same story twice (Hasson et al., 2009): if MD activity tracks the story in an idiosyncratic manner across individuals, then it should still be similar across two instances of the same story within an individual; however, if MD activity does not track the story, then it should not exhibit reliable time-courses even within an individual.

Therefore, we scanned several participants listening to stories twice, and then computed within-subject correlations (WSCs) for each network across the two instances. One group of participants (n=7) heard the stories repeatedly within the same scanning session (approximately one hour apart); another group (n=8) heard the stories in two sessions that were 6.5-21.5 months apart. These two groups did not differ from each other in their network WSCs, so their data were combined. In line with our findings above, WSCs in the LH language network (r=0.160) were stronger than in the RH language network (r=0.129, p<0.001), the MDfp network (r=0.083, $p\approx 0$) and the MDco network (r=0.097, $p<10^{-8}$). WSCs in the RH language network were stronger than those in the MDfp network ($p<10^{-4}$) and the MDco network (p=0.012), but the two latter networks did not differ (Figure 2.4A).

These WSCs are lower than the ISCs reported above; this effect was expected because WSCs are measured by correlating noisy signals from two single trials, whereas ISCs are measured by correlating a signal from one participant with an average (i.e., noise reduced) signal across all other participants. To better compare WSCs and ISCs, we thus re-computed ISCs by correlating signal time-courses across pairs of individual participants (Figure 2.4B). Now, ISCs appeared weaker than WSCs (i.e., signals across participants were less similar than signals within a participant), but both measures patterned similarly in terms of between-network differences (for all comparisons between WSCs and ISCs, $p>0.52$). Therefore, even across story repetitions within a given individual, MD network activity is significantly less reliable than language network activity, indicating that the former, but not the latter, tracks linguistic input closely.

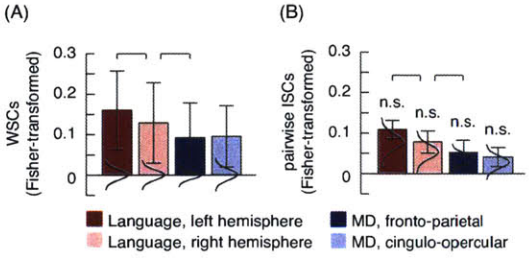

Figure 2.4. WSCs (A) and pairwise-ISCs (B) during story comprehension in the language and MD networks. Same conventions as in Figure 2.3. Note that pairwise ISCs are not significantly greater than expected by chance; however, they still exhibit differences across functional networks, similarly to WSCs.

2.2. Discussion

During story comprehension, a robust and reliable difference in neural activity distinguished between the language network and the MD network. The language network, particularly in the LH, showed relatively little individual differences in activity (high ISCs) due to close tracking of the story (high WSCs). In contrast, MD network activity was more idiosyncratic across individuals (low ISCs) with only weak tracking of the story (low WSCs). These findings suggest a novel typology of mental processes contributing to language comprehension: some processes implemented in the language network are stimulus-related and consistent across individuals; other processes, implemented in the MD network, are less tightly coupled to the input and appear more idiosyncratic. This distinction importantly constrains cognitive models of language processing.

Critically, characterizing the respective contributions of the language and MD networks to comprehension was methodologically possible due to localization of these networks using functional contrasts, individually for each participant. First, identifying networks functionally allows us to tie our findings to a wealth of prior literature characterizing the response profiles of those networks. Second, our approach takes into account inter-individual variability in the mapping of function onto anatomy by comparing functional regions across participants even when those regions do not align well spatially. Such variability, evident in the temporal cortex (Jones and Powell, 1970; Gloor, 1997; Wise et al., 2001) and especially in the frontal cortex (Amunts et al., 1999; Chein et al., 2002) where language and MD regions lie side by side

(Fedorenko et al., 2012a), renders anatomical localization precarious (Poldrack, 2006; Fischl et al., 2008; Frost and Goebel, 2012; Tahmasebi et al., 2012).

Indeed, pioneering studies of inter-subject correlations during language processing (Wilson et al., 2008; Lerner et al., 2011; Honey et al., 2012) computed ISCs for anatomical locations, assuming that the same location had a common function across participants. These studies revealed that broad cortical swaths show significant ISCs during comprehension, proposing a neural correlate of "shared understanding" across individuals (Hasson et al., 2012) yet offering no principled way to relate those regions to known functional divisions in the cortex. This issue was further complicated because studies had not directly contrasted regions to each other, and had usually reported only p-values but not the sizes of the correlations. By augmenting the ISC methodology with a single-participant functional localization approach, the present study provides one key characterization of the functional topography of ISCs, distinguishing between language and MD networks.

Within this topography, the role of MD regions in language comprehension is particularly interesting. Whereas task-based studies have demonstrated that MD regions scale their activity with increasing comprehension difficulty in numerous contexts (Stromswold et al., 1996; Stowe et al., 1998; Caplan et al., 1999; Fiez et al., 1999; Fiebach et al., 2002; Chee et al., 2003; Constable et al., 2004; Rodd et al., 2005; Chen et al., 2006; Nakic et al., 2006; Nieuwland et al., 2007; Novais-Santos et al., 2007; Hauk et al., 2008; Yarkoni et al., 2008; Carreiras et al., 2009; January et al., 2009; Peelle et al., 2009; Ye and Zhou, 2009; Barde et al., 2012; McMillan et al., 2012; McMillan et al., 2013), we demonstrate that they track natural language remarkably weakly even when it includes frequent occurrences of challenging linguistic features. Reconciling our data with past findings is thus challenging. Moreover, prior evidence suggests that MD regions track other naturalistic stimuli, such as audiovisual movies, with experiential features like "suspense" modulating MD activity similarly across individuals (Naci et al., 2014), possibly by influencing the frequency of attentional disengagement (Nakano et al., 2013). Does the domain-general MD network play a different role in language comprehension compared to its role in processing other naturalistic stimuli?

Perhaps MD regions are biased towards visual information (or audio-visual integration) in movies compared to the auditory information of stories (Michalka et al., 2015). Alternatively, MD regions may track both movies and stories, but fluctuations in MD activity during movie viewing could simply be slower, and thus more reliably measured, compared to the fast fluctuations during story comprehension. Therefore, evidence of stimulus tracking by MD regions during story comprehension might only be evident at high frequencies that cannot be measured with the temporally slow BOLD signal of fMRI. Finally, activity in MD regions may reflect internal fluctuations in domain-general attention or "focus" (Norman and Shallice, 1986; Chun et al., 2011) that may co-vary with the emotional manipulations in movies (Williams et al., 2016) but be relatively independent of input processing difficulty during natural language comprehension. This account is also consistent with previous findings of greater MD activity with increased linguistic demands in experimentally designed tasks, insofar as such tasks control the focus of participants more explicitly than naturalistic stories.

2.2.1 Conclusion

Using a combination of task-based functional localization in individual participants and a naturalistic cognition paradigm for comparing brain activity across participants, we characterize distinct contributions of the language network and MD network to story comprehension. Whereas activity in the language network is similar across individuals and closely tracks stories, activity in the MD network is more idiosyncratic and does not track linguistic input closely. These findings thus suggest a novel distinction between different mechanisms that underlie language processing based on individual differences in their processing patterns and their coupling to the linguistic input.

2.3. Materials and methods

2.3.1. Participants

Forty-five participants (30 females) between the ages of 18 and 50, recruited from the MIT student body and the surrounding community, were paid for participation. All participants were native English speakers, had normal hearing, and gave informed consent in accordance with the requirements of MIT's Committee on the Use of Humans as Experimental Subjects (COUHES).

2.3.2. Data acquisition

Structural and functional data were collected on a whole-body 3 Tesla Siemens Trio scanner with a 32-channel head coil at the Athinoula A. Martinos Imaging Center at the McGovern Institute for Brain Research at MIT. T1-weighted structural images were collected in 176 axial slices with 1mm isotropic voxels (repetition time (TR) = 2,530ms; echo time (TE) = 3.48ms). Functional, blood oxygenation level-dependent (BOLD) data were acquired using an EPI sequence with a 90° flip angle and using GRAPPA with an acceleration factor of 2; the following parameters were used: thirty-one 4.4mm thick near-axial slices acquired in an interleaved order (with 10% distance factor), with an in-plane resolution of 2.1mm × 2.1mm, FoV in the phase encoding (A >> P) direction 200mm and matrix size 96mm × 96mm, TR = 2000ms and TE = 30ms. The first 10s of each run were excluded to allow for steady state magnetization.

2.3.3. Design, materials and procedure

2.3.3.1. Language localizer task

Regions in the language network were localized using a reading task contrasting sentences and lists of unconnected, pronounceable nonwords (Fedorenko et al., 2010), in a standard blocked design with a counterbalanced order across runs (for timing parameters, see Table 2.1). Stimuli were presented one word/nonword at a time (Figure 2.1). For the first ten participants only, each trial ended with a memory probe and they had to indicate, via a button press, whether or not that probe had appeared in the preceding sequence of words/nonwords. For half of these participants, the localizer included an additional condition of unconnected word lists, for purposes of another experiment. The remaining 35 participants instead read the materials passively (we included a button-pressing event at the end of each trial, to help these participants remain alert and focused). Note that in the former version nonwords are more engaging than sentences because their memorization is harder, whereas in the latter version sentences were more engaging than nonwords because they have meaning. Importantly, this localizer has been shown to generalize across such manipulations, as the language network robustly and reliably shows a *sentences > nonwords* effect regardless of the task (Fedorenko et al., 2010). This localizer also generalizes across both visual and auditory presentations (Braze et al., 2011; Vagharchakian et al., 2012; Fedorenko, 2014).

2.3.3.2. MD localizer task

Regions in the MD network were localized with a spatial working memory game (Fedorenko et al., 2011) contrasting a hard version with an easy version. On each trial (8s), participants saw a 3×4 grid and kept track of eight (hard version) or four (easy version) locations that were sequentially flashed two at a time or one at a time, respectively (1s per flash, 4s total). Then, participants indicated their memory for these locations in a 2-alternative, forced-choice (2AFC) paradigm via a button press (3s total). Feedback was immediately provided upon choice (or lack thereof) (250ms). Trials began and ended with brief fixations (500ms and 250ms, respectively). Hard and easy conditions were presented in a standard blocked design (4 trials in a 32s block, 6 blocks per condition per run) with a counterbalanced order across runs. Each run included 4 blocks of fixation (16s each) and lasted a total of 448s. Thirty-nine participants completed 1-2 runs of the localizer. The remaining participants either provided poor-quality data (5 participants) or were not run on this task (1 participant). For this latter group, MD regions were localized with data from the language localizer task, namely, the *nonwords* > *sentences* contrast. Both the *hard* > *easy* contrast and the *nonwords* > *sentences* contrast have been previously demonstrated to robustly and reliably identify the MD network (Fedorenko et al., 2013).

Table 2.1. Timing parameters for the different versions of the language localizer task.

	Version		
	A	B	C
Number of participants	35	5	5
Task: Passive Reading or Memory?	PR	M	M
Words/nonwords per trial	12	12	12
Trial duration (ms)	6,000	6,000	6,000
Fixation	100	---	---
Presentation of each word / nonword	450	350	350
Fixation	500	300	300
Memory probe	---	1,000	1,000
Fixation	---	500	500
Trials per block	3	3	3
Block duration (s)	18	18	18
Blocks per condition (per run)	8	8	6
Conditions	Sentences Nonwords	Sentences Nonwords	Sentences Nonwords Word-lists
Fixation block duration (s)	14	18	18
Number of fixation blocks	5	5	4
Total run time (s)	358	378	396
Number of runs	2	2	2-3

2.3.3. Story comprehension task

Each subject listened to 1-4 stories over scanner-safe headphones (Sensimetrics, Malden, MA). In the main experiment and the first replication, stories were constructed based on publicly available fairy tales and short stories. These stories were edited to include a variety of linguistic phenomena that have been shown to increase processing difficulty and recruit the MD network, but do not occur with sufficiently high frequency in natural texts (for a sample text, see Appendix 2A). In the second replication, participants listened to an autobiographical story ("Pie-man," told by Jim O'Grady) recorded at a live storytelling event ("The Moth" storytelling event, NYC). Stories started and ended with 16s seconds of music and/or silence that were not analyzed.

After each story, participants answered 6-12 comprehension questions that required attentive listening (i.e., could not have been answered correctly based on common sense). For the main experiment and the first replication, participants answered 2AFC questions via a button press while in the scanner. For the second replication, participants filled in a 4AFC questionnaire once they got outside the scanner. For eight participants, answers to these questions were not collected. The remaining 37 participants demonstrated very good comprehension of the stories, with a negatively skewed accuracy distribution (mode=100%, median=87.5%, semi-interquartile range=12.85%).

2.3.4. Data preprocessing

Data preprocessing was carried out with SPM5 and custom MATLAB scripts. Preprocessing of anatomical data included normalization into a common space (Montreal Neurological Institute (MNI) template), resampling into 2mm isotropic voxels, and segmentation into probabilistic maps of the gray matter, white matter (WM) and cerebrospinal fluid (CSF). Preprocessing of functional data included motion correction, normalization, resampling into 2mm isotropic voxels, smoothing with a 4mm FWHM Gaussian filter and high-pass filtering at 200s.

Additional temporal preprocessing of data from the story comprehension runs was carried out using the CONN toolbox (Whitfield-Gabrieli and Nieto-Castanon, 2012) with default parameters, unless specified otherwise. Five temporal principal components of the BOLD signal time-courses extracted from the WM were regressed out of each voxel's time-course; signal originating in the CSF was similarly regressed out. Six principal components of the six motion parameters estimated during offline motion correction were also regressed out, as well as their first time derivative. Next, the residual signal was bandpass filtered (0.008–0.09 Hz) to preserve only low-frequency signal fluctuations (Cordes et al., 2001).

2.3.5. Data analysis

2.3.5.1. Modeling localizer data

For each localizer task, a general linear model estimated the effect size of each condition in each experimental run. These effects were each modeled with a boxcar function (representing entire blocks) convolved with the canonical hemodynamic response function (HRF). The model also included first-order temporal derivatives of these effects, as well as nuisance regressors representing entire experimental runs and offline-estimated motion parameters. The obtained beta weights were then used to compute the functional contrast of interest: *sentences* > *nonwords* for the language localizer, and *hard* > *easy* for the MD localizer.

2.3.5.2. Defining participant-specific language and MD regions

Language and MD regions were defined based on functional contrast maps from the localizer experiments. These maps were first restricted to include only gray matter voxels by excluding voxels that were more likely to belong to either the WM or the CSF based on SPM's probabilistic segmentation of the participant's structural data.

Then, regions in the language network were defined using group-constrained, participant-specific localization (Fedorenko et al., 2010). For each participant, the map of the *sentences* > *nonwords* contrast was intersected with binary masks that constrained the participant-specific language network to fall within areas where activations for this contrast are relatively likely across the population. These masks are based on a group-level representation of the contrast obtained from a previous sample. We used 8 such masks in the LH, including regions in the posterior, mid-posterior, mid-anterior and anterior temporal lobe, the angular gyrus, the middle frontal gyrus, the inferior frontal gyrus and its orbital part. These masks were mirror-projected onto the RH to create 8 homologous masks (the masks cover significant parts of the cortex, so their mirrored version is likely to encompass the RH homologue of the LH language network, despite possible hemispheric asymmetries in their precise locations). In each of the resulting 16 masks, a participant-specific language region was defined as the top 10% of voxels with the highest contrast values. This top *n%* approach ensures that functional regions can be defined in every participant and that their sizes are the same across participants, allowing for generalizable results (Nieto-Castañón and Fedorenko, 2012).

Regions in the MD network were similarly defined based on the *hard* > *easy* contrast in the spatial working memory game. Here, instead of using binary masks based on group-level data, we used anatomical masks (Tzourio-Mazoyer et al., 2002; for a similar approach, see Fedorenko et al., 2013; Blank et al., 2014). Nine masks were used in each hemisphere, including regions in the middle frontal gyrus and its orbital part, the opercular part of the inferior frontal gyrus, the precental gyrus, the posterior and inferior parts of the partieal lobe, the insula, and supplementary motor area and the cingulate cortex. The first five masks constitute the fronto-parietal MD sub-network, and the last three constitute the cingulo-opercular sub-network.

2.3.5.3. ISCs and WSCs

For each participant and functional region, BOLD signal time-courses recorded during story comprehension were extracted from each voxel beginning 6 seconds following the onset of the story (to exclude an initial rise in the hemodynamic response relative to fixation, which could increase ISCs). These time-courses we first temporally z-scored in each voxel and then averaged across voxels. Next, those signals were regressed against signals extracted from low-level auditory regions (defined anatomically around the postero-medial and antero-lateral sections of Heschl's gyrus bilaterally). Finally, for each participant and region, we computed Pearson's moment correlation coefficient between the residual time-course and the corresponding average residual time-course across the remaining participants (Lerner et al., 2011).

For each participant who listened to the same story on two occasions, we correlated the residual time-course in each region across the two trials. Because these WSCs are based on two single-trial signals, we also re-computed ISCs in a comparable manner; namely, for each participant and region, we correlated the residual time-course with the corresponding, individual residual time-course of each of the other participants, and averaged the resulting values.

ISCs/WSCs were Fisher-transformed prior to averaging and statistical testing in order to improve normality.

2.3.5.4 Statistical tests

In each region, ISCs/WSCs were tested for significance against an empirical null distribution based on 1,000 simulated signal time-courses that were generated by phase-randomization of the original data (Theiler et al., 1992). Individual distributions were each fit with a Gaussian and the resulting parameters were analytically combined across participants. The original ISCs/WSCs, also averaged across participants, were then z-scored relative to these parameters and converted to one-tailed p-values.

ISCs/WSCs were compared across networks using a linear, mixed-effects regression (Barr et al., 2013) implemented by the "lme4" package in R. In each experiment, ISCs/WSCs across all brain regions, participants and stories were modeled with a fixed effect of region and random intercepts for participant and story. The fixed effect estimates were combined across regions within each functional network (LH language, RH language, MDfp and MDco) and were pairwise compared to each other using the "multcomp" package in R. Hypotheses were two-tailed for the first experiment and one-tailed afterwards.

Statistical tests on WSC data were run on a sample including both participants who listened to the same story twice within the same scanning session and those who listened to the same story across two sessions. Prior to these analyses, we tested whether WSCs in the within-session and across-session datasets differed from each other. To this end, we performed a linear, mixed-effects regression analysis that modeled WSCs with a fixed effect of the interaction between brain region and dataset, random intercepts for participant and story, and a random slope for dataset varying by participant (this model was chosen because a fuller model failed to converge). Pairwise contrasts tested whether WSCs in each network were stronger across sessions than within a session. A similar approach was used for comparing WSCs to pairwise-ISCs. Here, contrasts tested whether pairwise differences between networks observed with WSCs were distinct from those observed with ISCs.

For all findings based on linear, mixed-effects regression analyses, similar results were obtained when ISCs/WSCs for each participant were first averaged across regions within each network and pairwise network comparisons (across participants) were then tested using exact permutation tests (Gill, 2007). Therefore, our results are independent of assumptions regarding data normality.

In each experiment, *p*-values are reported following false discovery rate (FDR) correction for multiple comparisons (Benjamini and Yekutieli, 2001).

2.4. Acknowledgements

We thank Alexander Paunov and Zach Mineroff for their help with data collection, Anastasia Vishnevetsky for her help with constructing the stories, Nancy Kanwisher and Ted Gibson for recording the stories, and Uri Hasson for providing the story for the second replication. We also thank the Athinoula A. Martinos Imaging Center at McGovern Institute for Brain Research at MIT, and the support team (Steve Shannon, Atsushi Takahashi, and Sheeba Arnold). Finally, we thank Nancy Kanwisher, Ted Gibson, John Duncan, and the audience at the CUNY Sentence Processing conference in San Diego for comments on earlier versions of this work. E.F. was supported by a K99/R00 award HD 057522 from NICHD.

References

Amunts K, Schleicher A, Bürgel U, Mohlberg H, Uylings H, Zilles K (1999) Broca's region revisited: Cytoarchitecture and intersubject variability. Journal of Comparative Neurology 412:319-341.

Barde LH, Yeatman JD, Lee ES, Glover G, Feldman HM (2012) Differences in neural activation between preterm and full term born adolescents on a sentence comprehension task: Implications for educational accommodations. Developmental Cognitive Neuroscience 2:S114-S128.

Barr DJ, Levy R, Scheepers C, Tily HJ (2013) Random effects structure for confirmatory hypothesis testing: Keep it maximal. Journal of Memory and Language 68:255-278.

Bates E, Wilson SM, Saygin AP, Dick F, Sereno MI, Knight RT, Dronkers NF (2003) Voxel-based lesion-symptom mapping. Nature Neuroscience 6:448-450.

Bautista A, Wilson SM (2016) Neural responses to grammatically and lexically degraded speech. Language, Cognition and Neuroscience 31:567-574.

Bedny M, McGill M, Thompson-Schill SL (2008) Semantic adaptation and competition during word comprehension. Cerebral Cortex 18:2574-2585.

Benjamini Y, Yekutieli D (2001) The control of the false discovery rate in multiple testing under dependency. Annals of statistics:1165-1188.

Binder JR, Frost JA, Hammeke TA, Cox RW, Rao SM, Prieto T (1997) Human brain language areas identified by functional magnetic resonance imaging. The Journal of Neuroscience 17:353-362.

Blank I, Kanwisher N, Fedorenko E (2014) A functional dissociation between language and multiple-demand systems revealed in patterns of BOLD signal fluctuations. Journal of Neurophysiology 112:1105-1118.

Blank I, Balewski Z, Mahowald K, Fedorenko E (2016) Syntactic processing is distributed across the language system. NeuroImage 127:307-323.

Bozic M, Fonteneau E, Su L, Marslen - Wilson WD (2015) Grammatical analysis as a distributed neurobiological function. Human Brain Mapping 36:1190-1201.

Braze D, Mencl WE, Tabor W, Pugh KR, Constable RT, Fulbright RK, Magnuson JS, Van Dyke JA, Shankweiler DP (2011) Unification of sentence processing via ear and eye: An fMRI study. Cortex 47:416-431.

Caplan D, Alpert N, Waters G (1999) PET studies of syntactic processing with auditory sentence presentation. NeuroImage 9:343-351.

Carreiras M, Riba J, Vergara M, Heldmann M, Münte TF (2009) Syllable congruency and word frequency effects on brain activation. Human Brain Mapping 30:3079-3088.

Chee MWL, Westphal C, Goh J, Graham S, Song AW (2003) Word frequency and subsequent memory effects studied using event-related fMRI. NeuroImage 20:1042-1051.

Chein JM, Fissell K, Jacobs S, Fiez JA (2002) Functional heterogeneity within Broca's area during verbal working memory. Physiology and Behavior 77:635-639.

Chen E, West WC, Waters G, Caplan D (2006) Determinants of BOLD signal correlates of processing object-extracted relative clauses. Cortex 42:591-604.

Chiarello C, Banich MT, Mack M (2003) Parallel systems for processing language: Hemispheric complementarity in the normal brain. In: Mind, Brain, and Language: Multidisciplinary Perspectives (Banich MT, Mack M, eds), pp 229-247. Mahwah, NJ: Lawrence Erlbaum.

Chun MM, Golomb JD, Turk-Browne NB (2011) A taxonomy of external and internal attention. Annual Review of Psychology 62:73-101.

Collins MJ (1996) A new statistical parser based on bigram lexical dependencies. In: Proceedings of the 34th Annual Meeting on Association for Computational Linguistics, pp 184-191. Stroudsburg, PA: Association for Computational Linguistics.

Constable RT, Pugh KR, Berroya E, Mencl WE, Westerveld M, Ni W, Shankweiler D (2004) Sentence complexity and input modality effects in sentence comprehension: an fMRI study. NeuroImage 22:11-21.

Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, Quigley MA, Meyerand ME (2001) Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. American Journal of Neuroradiology 22:1326-1333.

D'Esposito M, Postle BR (2015) The cognitive neuroscience of working memory. Annual Review of Psychology 66:115-142.

Damasio AR (1992) Aphasia. New England Journal of Medicine 326:531-539.

Dosenbach NUF, Visscher KM, Palmer ED, Miezin FM, Wenger KK, Kang HC, Burgund ED, Grimes AL, Schlaggar BL, Petersen SE (2006) A core system for the implementation of task sets. Neuron 50:799-812.

Dosenbach NUF, Fair DA, Miezin FM, Cohen AL, Wenger KK, Dosenbach RAT, Fox MD, Snyder AZ, Vincent JL, Raichle ME (2007) Distinct brain networks for adaptive and stable task control in humans. Proceedings of the National Academy of Sciences 104:11073-11078.

Duncan J (2010) The multiple-demand (MD) system of the primate brain: Mental programs for intelligent behaviour. Trends in Cognitive Sciences 14:172-179.

Fedorenko E (2014) The role of domain-general cognitive control in language comprehension. Frontiers in Psychology 5:335.

Fedorenko E, Thompson-Schill SL (2014) Reworking the language network. Trends in Cognitive Sciences 18:120-126.

Fedorenko E, Behr MK, Kanwisher N (2011) Functional specificity for high-level linguistic processing in the human brain. Proceedings of the National Academy of Sciences 108:16428-16433.

Fedorenko E, Duncan J, Kanwisher N (2012a) Language-selective and domain-general regions lie side by side within Broca's area. Current Biology 22:2059-2062.

Fedorenko E, Nieto-Castañón A, Kanwisher N (2012b) Lexical and syntactic representations in the brain: an fMRI investigation with multi-voxel pattern analyses. Neuropsychologia 50:499-513.

Fedorenko E, Duncan J, Kanwisher N (2013) Broad domain generality in focal regions of frontal and parietal cortex. Proceedings of the National Academy of Sciences 110:16616-16621.

Fedorenko E, Hsieh P-J, Nieto-Castañón A, Whitfield-Gabrieli S, Kanwisher N (2010) New method for fMRI investigations of language: Defining ROIs functionally in individual subjects. Journal of Neurophysiology 104:1177-1194.

Fiebach CJ, Friederici AD, Müller K, Von Cramon DY (2002) fMRI evidence for dual routes to the mental lexicon in visual word recognition. Journal of Cognitive Neuroscience 14:11-23.

Fiez JA, Balota DA, Raichle ME, Petersen SE (1999) Effects of lexicality, frequency, and spelling-to-sound consistency on the functional anatomy of reading. Neuron 24:205-218.

Fischl B, Rajendran N, Busa E, Augustinack J, Hinds O, Yeo BTT, Mohlberg H, Amunts K, Zilles K (2008) Cortical folding patterns and predicting cytoarchitecture. Cerebral Cortex 18:1973-1980.

Frost MA, Goebel R (2012) Measuring structural–functional correspondence: Spatial variability of specialised brain regions after macro-anatomical alignment. NeuroImage 59:1369-1381.

Futrell R, Mahowald K, Gibson E (2015) Large-scale evidence of dependency length minimization in 37 languages. Proceedings of the National Academy of Sciences 112:10336-10341.

George MS, Parekh PI, Rosinsky N, Ketter TA, Kimbrell TA, Heilman KM, Herscovitch P, Post RM (1996) Understanding emotional prosody activates right hemisphere regions. Archives of Neurology 53:665-670.

Geschwind N (1970) The organization of language and the brain. In: Science: Citeseer.

Gill PMW (2007) Efficient calculation of p-values in linear-statistic permutation significance tests. Journal of Statistical Computation and Simulation 77:55-61.

Gloor P (1997) The temporal lobe and limbic system. New York, NY: Oxford University Press.

Hasson U, Malach R, Heeger DJ (2010) Reliability of cortical activity during natural stimulation. Trends in Cognitive Sciences 14:40-48.

Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R (2004) Intersubject synchronization of cortical activity during natural vision. science 303:1634-1640.

Hasson U, Ghazanfar AA, Galantucci B, Garrod S, Keysers C (2012) Brain-to-brain coupling: A mechanism for creating and sharing a social world. Trends in Cognitive Sciences 16:114-121.

Hasson U, Avidan G, Gelbard H, Vallines I, Harel M, Minshew N, Behrmann M (2009) Shared and idiosyncratic cortical activation patterns in autism revealed under continuous real - life viewing conditions. Autism Research 2:220-231.

Hauk O, Davis MH, Pulvermüller F (2008) Modulation of brain activity by multiple lexical and word form variables in visual word recognition: A parametric fMRI study. NeuroImage 42:1185-1195.

Hindy NC, Altmann GTM, Kalenik E, Thompson-Schill SL (2012) The effect of object state-changes on event processing: Do objects compete with themselves? The Journal of Neuroscience 32:5795-5803.

Hindy NC, Solomon SH, Altmann GTM, Thompson-Schill SL (2015) A cortical network for the encoding of object change. Cerebral Cortex 25:884-894.

Honey CJ, Thompson CR, Lerner Y, Hasson U (2012) Not lost in translation: Neural responses shared across languages. The Journal of Neuroscience 32:15277-15283.

January D, Trueswell JC, Thompson-Schill SL (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 21:2434-2444.

Jones EG, Powell TP (1970) An anatomical study of converging sensory pathways within the cerebral cortex of the monkey. Brain 93:793-820.

Jung-Beeman M (2005) Bilateral brain processes for comprehending natural language. Trends in Cognitive Sciences 9:512-518.

Keller TA, Carpenter PA, Just MA (2001) The neural bases of sentence comprehension: a fMRI examination of syntactic and lexical processing. Cerebral Cortex 11:223-237.

Lerner Y, Honey CJ, Silbert LJ, Hasson U (2011) Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. The Journal of Neuroscience 31:2906-2915.

Mantini D, Corbetta M, Romani GL, Orban GA, Vanduffel W (2013) Evolutionarily novel functional networks in the human brain? The Journal of Neuroscience 33:3259-3275.

McMillan CT, Clark R, Gunawardena D, Ryant N, Grossman M (2012) fMRI evidence for strategic decision-making during resolution of pronoun reference. Neuropsychologia 50:674-687.

McMillan CT, Coleman D, Clark R, Liang T-W, Gross RG, Grossman M (2013) Converging evidence for the processing costs associated with ambiguous quantifier comprehension. Frontiers in psychology 4:153.

Menenti L, Gierhan SME, Segaert K, Hagoort P (2011) Shared language overlap and segregation of the neuronal infrastructure for speaking and listening revealed by functional MRI. Psychological Science 22:1173-1182.

Mesulam MM, Thompson CK, Weintraub S, Rogalski EJ (2015) The Wernicke conundrum and the anatomy of language comprehension in primary progressive aphasia. Brain 138:2423-2437.

Michalka SW, Kong L, Rosen ML, Shinn-Cunningham BG, Somers DC (2015) Short-term memory for space and time flexibly recruit complementary sensory-biased frontal lobe attention networks. Neuron 87:882-892.

Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. Annual Review of Neuroscience 24:167-202.

Mirman D, Chen Q, Zhang Y, Wang Z, Faseyitan OK, Coslett HB, Schwartz MF (2015) Neural organization of spoken language revealed by lesion-symptom mapping. Nature Communications 6.

Monti MM, Osherson DN (2012) Logic, language and the brain. Brain Research 1428:33-42.

Monti MM, Parsons LM, Osherson DN (2009) The boundaries of language and thought in deductive inference. Proceedings of the National Academy of Sciences 106:12554-12559.

Monti MM, Parsons LM, Osherson DN (2012) Thought beyond language: neural dissociation of algebra and natural language. Psychological Science 23:914-922.

Naci L, Cusack R, Anello M, Owen AM (2014) A common neural code for similar conscious experiences in different individuals. Proceedings of the National Academy of Sciences 111:14277-14282.

Nakano T, Kato M, Morito Y, Itoi S, Kitazawa S (2013) Blink-related momentary activation of the default mode network while viewing videos. Proceedings of the National Academy of Sciences 110:702-706.

Nakic M, Smith BW, Busis S, Vythilingam M, Blair RJR (2006) The impact of affect and frequency on lexical decision: The role of the amygdala and inferior frontal cortex. NeuroImage 31:1752-1761.

Nieto-Castañón A, Fedorenko E (2012) Subject-specific functional localizers increase sensitivity and functional resolution of multi-subject analyses. NeuroImage 63:1646-1669.

Nieuwland MS, Petersson KM, Van Berkum JJA (2007) On sense and reference: Examining the functional neuroanatomy of referential processing. NeuroImage 37:993-1004.

Nomura EM, Gratton C, Visser RM, Kayser A, Perez F, D'Esposito M (2010) Double dissociation of two cognitive control networks in patients with focal brain lesions. Proceedings of the National Academy of Sciences 107:12017-12022.

Norman DA, Shallice T (1986) Attention to action. In: Consciousness and Self-Regulation, pp 1-18. New York, NY: Springer.

Novais-Santos S, Gee J, Shah M, Troiani V, Work M, Grossman M (2007) Resolving sentence ambiguity with planning and working memory resources: Evidence from fMRI. NeuroImage 37:361-378.

Peelle JE, Troiani V, Wingfield A, Grossman M (2009) Neural processing during older adults' comprehension of spoken sentences: Age differences in resource allocation and connectivity. Cerebral Cortex 20:bhp142.

Poldrack RA (2006) Can cognitive processes be inferred from neuroimaging data? Trends in Cognitive Sciences 10:59-63.

Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL (2011) Functional network organization of the human brain. Neuron 72:665-678.

Price CJ (2012) A review and synthesis of the first 20 years of PET and fMRI studies of heard speech, spoken language and reading. NeuroImage 62:816-847.

Rodd JM, Davis MH, Johnsrude IS (2005) The neural mechanisms of speech comprehension: fMRI studies of semantic ambiguity. Cerebral Cortex 15:1261-1269.

Roland D, Dick F, Elman JL (2007) Frequency of basic English grammatical structures: A corpus analysis. Journal of Memory and Language 57:348-379.

Sreenivasan KK, Curtis CE, D'Esposito M (2014) Revisiting the role of persistent neural activity during working memory. Trends in Cognitive Sciences 18:82-89.

Stowe LA, Broere CAJ, Paans AMJ, Wijers AA, Mulder G, Vaalburg W, Zwarts F (1998) Localizing components of a complex task: Sentence processing and working memory. Neuroreport 9:2995-2999.

Stromswold K, Caplan D, Alpert N, Rauch S (1996) Localization of syntactic comprehension by positron emission tomography. Brain and language 52:452-473.

Tahmasebi AM, Davis MH, Wild CJ, Rodd JM, Hakyemez H, Abolmaesumi P, Johnsrude IS (2012) Is the link between anatomical structure and function equally strong at all cognitive levels of processing? Cerebral Cortex 22:1593-1603.

Theiler J, Eubank S, Longtin A, Galdrikian B, Farmer JD (1992) Testing for nonlinearity in time series: The method of surrogate data. Physica D: Nonlinear Phenomena 58:77-94.

Thompson-Schill SL, Bedny M, Goldberg RF (2005) The frontal lobes and the regulation of mental activity. Current Opinion in Neurobiology 15:219-224.

Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M (2002) Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15:273-289.

Vagharchakian L, Dehaene-Lambertz G, Pallier C, Dehaene S (2012) A temporal bottleneck in the language comprehension network. The Journal of Neuroscience 32:9089-9102.

Whitfield-Gabrieli S, Nieto-Castanon A (2012) Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. Brain connectivity 2:125-141.

Williams J, Stönner C, Wicker J, Krauter N, Derstroff B, Bourtsoukidis E, Klüpfel T, Kramer S (2016) Cinema audiences reproducibly vary the chemical composition of air during films, by broadcasting scene specific emissions on breath. Scientific reports 6:25464.

Wilson SM, Molnar-Szakacs I, Iacoboni M (2008) Beyond superior temporal cortex: Intersubject correlations in narrative speech comprehension. Cerebral Cortex 18:230-242.

Wise RJS, Scott SK, Blank SC, Mummery CJ, Murphy K, Warburton EA (2001) Separate neural subsystems within "Wernicke's area". Brain 124:83-95.

Yang J (2014) The role of the right hemisphere in metaphor comprehension: A meta - analysis of functional magnetic resonance imaging studies. Human Brain Mapping 35:107-122.

Yarkoni T, Speer NK, Balota DA, McAvoy MP, Zacks JM (2008) Pictures of a thousand words: Investigating the neural mechanisms of reading with extremely rapid event-related fMRI. NeuroImage 42:973-987.

Ye Z, Zhou X (2009) Conflict control during sentence comprehension: fMRI evidence. NeuroImage 48:280-290.

Appendix 2A: A sample story and comprehension questions

At ten years old, I could not figure out what it was that this Elvis Presley guy had that the rest of us boys did not have. He seemed to be no different from the rest of us. He was simply a man who had a head, two arms and two legs. It must have been something pretty superlative that he had hidden away, because he had every young girl at the orphanage wrapped around his little finger.

At about nine o'clock on Saturday morning, I figured a good solution was to ask Eugene Correthers, who was one of the older and smarter boys, what it was that made this Elvis guy so special. He told me that it was not anything about Elvis's personality, but his wavy hair, and the way he moved his body. About a half an hour later, the boys in the orphanage called down to the main dining room by the matron were told that they were all going to downtown Jacksonville, Florida to get a new pair of Buster Brown shoes and a haircut. That is when I got this big idea, which hit me like a ton of bricks. If the Elvis haircut was the big secret, then Elvis's haircut I was going to get.

I was going to have my day in the sun, and all the way to town that was all I talked about. The fact that I was getting an Elvis haircut, not just the simple fact that we were getting out of the orphanage, made me particularly loquacious. I told everybody, including the orphanage matron I normally feared, that I was going to look just like Elvis Presley and that I would learn to move around just like he did and that I would be rich and famous one day, just like him. The matron understood my idea was something that I was really excited about and said nothing.

When I got my new Buster Brown shoes, I was smiling from ear to ear. Those shoes, they shined really brightly, and I liked looking at the bones in my feet, which I had never seen before, through a special x-ray machine they had in the shoe store that made the bones in your feet look green. I was now almost ready to go back to the orphanage and practice being like the man who all the girls loved, since I had my new Buster Brown shoes. It was the new haircut, though, that I needed to complete my new look.

We finally arrived at the unassuming, unembellished barbershop, where they cut our hair for free because we were orphans. Even though we were supposed to slowly wait to be called, I ran straight up to one of the barber chairs and climbed up onto the board the barber placed across the arms to make me sit up higher. I looked at the man and said, with a beaming smile on my face, "I want an Elvis haircut. Can you make my hair like Elvis's?" I asked. The barber, who was a genial young man, grinned back at me and said that he would try his best.

I was so happy when he started to cut my hair, but just as he started to cut, the matron, who had been watching me and had a look as cold as ice, motioned for him to come over to where she was standing. She whispered something into his ear that caused the barber to shake his head, like he was telling her, "No". In response, the matron walked over to a little man sitting in an office chair that squeaked as it rolled around the floor and spoke to him. It was the little man who then walked over and said something to the man who was cutting my hair. The next thing I knew, the man who was cutting my hair told me that he was no longer allowed to give me an Elvis cut.

"Why not?" I cried desperately.

The kindly barber stopped by the matron did not answer, but from his expression, I could tell that he wished he could cut it as I had asked.

Within a few minutes, it wasn't an Elvis haircut, but a short buzz cut that the barber had given me. When he finished shaving off all my hair and made me smell real good with his powder, the barber handed me a nickel and told me to go outside to the snack machine and buy myself a candy bar. I handed him the nickel back and told him that I was not hungry. "I'm so sorry, baby," he said, as I climbed out of his barber chair. "I am not a baby," I said, as I wiped the tears from my eyes.

I then sat down on the floor and brushed away the hair that had accumulated on my shiny new Buster Brown shoes. My head was no longer in the clouds, and I got up off the floor, brushed off my short pants, and walked sullenly towards the door.

The matron was smiling at me sort of funny like.

The barber upset by the matron said to her, "You are just a damn bitch, lady."

She yelled back at him at the top of her lungs, before walking toward the office, as fast as she could.

To show his anger, the man hit the wall with his hand and then walked outside where he stood against the brick wall, smoking a cigarette. I understood right there my haircut was something that had been out of the power of the barber and then I slowly walked outside to join the man. He looked down, smiled at me, then he patted me on the top of my bald as a coot head. It was a fact of my life that I was not gonna have hair that was anything like Elvis's anytime soon. I then looked up at the barber with my wet red eyes and asked, "Do you know if Elvis Presley has green bones?"

1. Why was the boy interested in Elvis?
    A. Girls at the orphanage liked Elvis
    B. Elvis had a lot of money

2. What made Elvis special, in the opinion of Eugene Correthers?
    A. Elvis's personality
    B. Wavy hair

3. On the bus, how did the boy behave?
    A. Talked a lot
    B. Was very quiet

4. What was the barber's initial reaction to the boy's request?
    A. Said he couldn't do it
    B. Said he'd try his best

5. What did the barber hand to the boy when he finished the haircut?
    A. A candy
    B. A nickel

6. What did the barber do to show his anger?
    A. Pushed the matron aside
    B. Hit the wall

Part B

Distributed processing throughout the high-level language network in the left hemisphere

[This page intentionally left blank]

Chapter 3

Syntactic processing is distributed throughout the language network

With Zuzanna Balewski, Kyle Mahowald and Evelina Fedorenko

Abstract

Language comprehension recruits an extended set of regions in the human brain. Is syntactic processing localized to a particular region or regions within this network, or is it distributed across the entire ensemble of brain regions that support high-level linguistic processing? Evidence from aphasic patients is more consistent with the latter possibility: damage to many different language regions and to white-matter tracts connecting them has been shown to lead to similar syntactic comprehension deficits. However, brain imaging investigations of syntactic processing continue to focus on particular regions within the language network, often parts of Broca's area and regions in the posterior temporal cortex. We hypothesized that, whereas the entire language network is in fact sensitive to syntactic complexity, the effects in some regions may be difficult to detect because of the overall lower response to language stimuli. Using an individual-participants approach to localizing the language network, shown in prior work to be more sensitive than traditional group analyses, we indeed find responses to syntactic complexity throughout this network, consistent with the findings from the neuropsychological patient literature. We speculate that such distributed nature of syntactic processing could perhaps imply that syntax is inseparable from other aspects of language comprehension (e.g., lexico-semantic processing), in line with current linguistic and psycholinguistic theories and evidence. Neuroimaging investigations of syntactic processing thus need to expand their scope to include the entire network of high-level language processing regions in order to fully understand how syntax is instantiated in the human brain.

3.1. Introduction

Language processing is supported by an extended network of brain regions, primarily in the left frontal and temporal lobes (Binder et al., 1997; Fedorenko et al., 2010). Whereas evidence from both the patient and neuroimaging literatures strongly suggests that this network is selectively engaged in linguistic processes and not in other cognitive processes (e.g., Dronkers et al., 1998; Varley et al., 2005; Fedorenko et al., 2011; Willems et al., 2011; Fedorenko et al., 2012a; Monti et al., 2012), the division of linguistic labor among its constituent regions is still heavily debated. A key question for understanding the internal structure of the language network is to what extent different aspects of language comprehension are localized to particular regions within the network versus distributed across the entire network. The answer to this question will reveal which functions are implemented in distinct neural circuits and which functions share neural resources. These organizational principles of neural architecture might, in turn, illuminate the cognitive architecture of the human language faculty (for similar inferences from neural to cognitive architectures in perception, see e.g., Kanwisher, 2010). Here, we specifically focus on syntactic processing: is it localized or distributed across the language network?

Prior literature addressing this issue provides conflicting evidence, such that neuropsychological evidence – on the whole – supports a distributed view of syntactic processing whereas neuroimaging evidence appears to support a more localized view. On the one hand, investigations of patients with brain damage have revealed that lesions to many different parts of the language network can cause similar syntactic comprehension difficulties. Such regions include Broca's region in the inferior frontal gyrus (Caramazza and Zurif, 1976; Schwartz et al., 1980; Caplan and Futter, 1986; Zurif et al., 1993; Grodzinsky, 2000), the arcuate fasciculus and/or the extreme capsule (Caramazza and Zurif, 1976; Papoutsi et al., 2011; Rolheiser et al., 2011; Tyler et al., 2011; Wilson et al., 2011), posterior temporal regions (Samuels and Benson, 1979; Selnes et al., 1983; Basso et al., 1985; Tramo et al., 1988; Caplan et al., 1996; Bastiaanse and Edwards, 2004; Wilson and Saygin, 2004; Amici et al., 2007; Tyler et al., 2011; Thothathiri et al., 2012), and anterior temporal regions (Dronkers et al., 1994; Dronkers et al., 2004; Magnusdottir et al., 2013). For instance, lesions in all of these regions can impair the interpretation of semantically reversible sentences, such as THE BOY CHASED THE GIRL, in which meaboning (who did what to whom) depends on syntactic form (i.e., word order, function words, and functional morphology). Consequently, some have argued that syntactic processing is supported by the language network as a whole (Caplan et al., 1996; Dick et al., 2001; Wilson and Saygin, 2004; Mesulam et al., 2015).

On the other hand, many neuroimaging studies employing syntactic manipulations have found activations not across the entire language network but, instead, restricted to a subset of the network, most commonly in the inferior frontal and posterior temporal regions (Just et al., 1996; Stromswold et al., 1996; Cooke et al., 2002; Ben-Shachar et al., 2003; Wartenburger et al., 2003; Constable et al., 2004; Bornkessel et al., 2005; Fiebach et al., 2005; Caplan et al., 2008; Meltzer et al., 2010; Peelle et al., 2010; see Friederici, 2011, for a recent meta-analysis). These studies suggest a localized view of syntactic processing, in line with many proposals that link syntax to Broca's area (Ullman, 2004; Bornkessel and Schlesewsky, 2006; Grodzinsky and Friederici, 2006; Grodzinsky and Santi, 2008; Friederici, 2009; Baggio and Hagoort, 2011; Friederici, 2011; Tyler et al., 2011; Friederici, 2012; Duffau et al., 2014).

How can we reconcile these two sets of conflicting findings? One possibility is that the localized activation patterns in neuroimaging studies result from (i) the use of group analyses, which suffer from sensitivity loss due to inter-individual variability in the precise locations of activation peaks (e.g., Nieto-Castañón and Fedorenko, 2012); and (ii) differences across brain regions in the overall strength of response to language stimuli. In highly language-responsive regions one might expect relatively wide neighborhoods of strong activation, so that overlap across participants could be evident despite inter-individual variability in peak location. In regions that are language-selective but respond only weakly to language stimuli, however, one might expect smaller and shallower activation neighborhoods surrounding the (low) peaks, so that overlapping activations across participants are less likely to emerge. Such reasoning suggests that neuroimaging methods that take into account inter-individual variability may be able to find evidence for distributed, rather than localized, syntactic processing. Therefore, here we use an individual-participants approach (Fedorenko et al., 2010) that allows us to narrow in on the high-level language processing regions in each individual brain. We measure the effect of syntactic complexity on the response of these individually localized regions and show that, in fact, syntactic complexity modulates neural responses throughout the language network, consistent with the evidence from the patient literature.

3.2. Materials and methods

To test for sensitivity to syntactic demands, we chose a commonly used syntactic complexity manipulation: the contrast between subject- and object-extracted relative clauses, as in (1) (see also Figure 3.1).

(1) a. Subject-extracted relative clause: THE STAR THAT IS GREETING THE CIRCLR

b. Object-extracted relative clause: THE CIRCLE THAT THE STAR IS GREETING

In both (1a) and (1b), the verb phrase IS GREETING has two arguments (i.e., dependents): a subject who is doing the greeting (THE STAR), and an object who is being greeted (THE CIRCLE). However, the two sentences critically differ in the distance separating the verb phrase from its two dependents. Specifically, in the subject-extracted relative clause (1a), the dependencies are local: both the word THAT (which refers to THE STAR) and the object THE CIRCLE connect locally to the verb phrase IS GREETING. In contrast, the object-extracted relative clause (1b) has a more complex dependency structure: the verb phrase IS GREETING is separated from its object, THE CIRCLE, by the subject THE STAR. An appealing feature of this contrast is that a variety of factors that have been shown to affect sentence comprehension (Tanenhaus and Trueswell, 1995; Gibson and Pearlmutter, 1998) are matched across the two conditions, including lexical-level factors (the words are identical) and plausibility. So, only the dependency structure (i.e., syntax) varies.

Across many languages, object-extracted relative clauses like (1b) have been shown to cause comprehension difficulty compared to subject-extracted relative clauses like (1a), as reflected in a variety of dependent measures including reading times and response accuracies to comprehension questions (e.g., English: Wanner & Maratsos, 1978; King & Just, 1991; Gibson, 1998; Grodner & Gibson, 2005; French: Holmes & O'Regan, 1981; Baudiffier, Caplan, Gaonac'h, & Chesnet, 2011; German: Mecklinger, Schriefers, Steinhauer, & Friederici, 1995; Schriefers, Friederici, & Kuhn, 1995; Dutch: Frazier, 1987; Mak, Vonk, & Schriefers, 2002; 2006; Japanese: Miyamoto & Nakamura 2003; Ishizuka, Nakatani, & Gibson, 2003; Ueno & Garnsey, 2008; Korean: O'Grady, Lee & Choo, 2003; Kwon, Polinsky, & Kluender, 2006; Kwon, Gordon, Kluender, & Polinsky, 2010; Russian: Levy, Fedorenko & Gibson, 2013). Therefore, the contrast between object- and subject-extracted relative clauses is considered by many to be a marker of syntactic processing, and has been used widely in both investigations of individuals with aphasia and brain imaging studies.

As mentioned above, in previous neuroimaging work, such contrasts between object- and subject-extractions as well as other, similar contrasts have produced activations largely restricted to Broca's area, the surrounding regions in the inferior frontal gyrus and the posterior parts of the middle (and sometimes superior) temporal gyrus. Other regions in the language network – such as the orbital portions of the inferior frontal gyrus or the anterior temporal regions – did not show reliable responses. However, this data pattern does not necessarily imply that the former regions are significantly more sensitive to the syntactic manipulation than the latter. Before such a claim is put forward, two methodological issues warrant consideration.

The first issue concerns the sensitivity of fMRI analysis methods to syntactic complexity effects. The vast majority of previous studies have relied on traditional group analyses, where individual brains are transformed into a common space and their contrast maps are then averaged across participants, assuming a shared mapping of functional regions onto anatomy. Although such methods can be effective in detecting large regions of activation that align well across individuals, they suffer from sensitivity loss due to inter-individual anatomical and functional variability (e.g., Saxe et al., 2006; Nieto-Castañón and Fedorenko, 2012), which has been shown to be especially pronounced in the frontal and temporal cortices (e.g., Frost and Goebel, 2012; Tahmasebi et al., 2012). As a result, even when every participant shows a robust response to syntactic complexity manipulations individually, the effect may get "washed out" by group averaging (see, Fedorenko and Kanwisher, 2011, for an example).

The second issue regards the validity of statistical tests. Namely, observing that some regions show a significant syntactic complexity effect, whereas others do not, cannot be taken as evidence that regions differ from one another in how engaged they are by syntactic processing. Such an inference would only be licensed by directly comparing contrast effects across regions, with some regions showing a stronger difference between the responses to object- and subject-extractions, compared to other regions. In other words, a region by extraction-type interaction is needed (see Nieuwenhuis et al., 2011 for a recent discussion).

In summary, in order to argue that only a particular subset of the language network is engaged in syntactic processing (or, more generally, that different parts of the language network support distinct computations), it is important to use methods that take into account inter-individual variability in the exact location of syntactic effects and explicitly test hypotheses of interest. One way to take inter-individual variability into account in the second-level analyses is by using a functional "localizer" contrast that narrows in on the relevant functional subset of each individual brain, in order to then examine the responses of those functionally defined regions to the critical conditions of interest. Thus, we here use a functional localizer for brain regions that support high-level linguistic processing (Fedorenko et al., 2010), which robustly activates the key language-responsive regions in the frontal, temporal and parietal cortices. We then employ a standard sentence-picture matching paradigm with object- and subject-extractions to examine whether syntactic complexity affects the response of these brain regions. In addition to testing the significance of this effect in each region, we also test for a region by extraction-type interaction, to assess whether some regions are more sensitive than others to syntactic complexity.

3.2.1. Participants

Thirteen participants (10 females) between the ages of 18 and 30 – students at MIT and members of the surrounding community – were paid for their participation. Participants were right-handed native speakers of English, naïve to the purposes of the study. All participants gave informed consent in accordance with the requirements of MIT's Committee On the Use of Humans as Experimental Participants (COUHES).

3.2.2. Data acquisition

Structural and functional data were collected on the whole-body 3 Tesla Siemens Trio scanner with a 32-channel head coil at the Athinoula A. Martinos Imaging Center at the McGovern Institute for Brain Research at MIT. T1-weighted structural images were collected in 128 axial slices with 1.33 mm isotropic voxels (TR=2000ms, TE=3.39ms). Functional, blood oxygenation level dependent (BOLD), data were acquired using an EPI sequence (with a 90° flip angle and using GRAPPA with an acceleration factor of 2), with the following acquisition parameters: thirty-one 4 mm thick near-axial slices acquired in the interleaved order (with 10% distance factor), 2.1mmx2.1mm in-plane resolution, FoV in the phase encoding (A >> P) direction 200mm and matrix size 96x96, TR=2000ms and TE=30ms. The first 10s of each run were excluded to allow for steady state magnetization.

3.2.3. Design, materials and procedure

Each participant performed the language localizer task (Fedorenko *et al.*, 2010) and the critical syntactic-processing task. Some participants also completed one or two additional tasks for unrelated studies. The entire scanning session lasted approximately 2 hours.

3.2.3.1. Language localizer task

Participants read sentences (e.g., A RUSTY LOCK WAS FOUND IN THE DRAWER) and lists of unconnected, pronounceable nonwords (e.g., DAP DRELLO SMOP UL PLID KAV CRE REPLODE) in a blocked design. Each stimulus consisted of eight words/nonwords. For details of how the language materials were constructed, see Fedorenko et al. (2010). The materials are available at http://web.mit.edu/evelina9/www/funcloc/funcloc_localizers.html. Stimuli were presented in the center of the screen, one word/nonword at a time, at the rate of 350ms per word/nonword. Each stimulus was followed by a 300ms blank screen, a memory probe (presented for 1350ms), and again a blank screen for 350ms, for a total trial duration of 4.8s. Participants were asked to decide whether the probe appeared in the preceding stimulus by pressing one of two buttons. (In previous work we established that similar activations obtain with a passive-reading task; see Fedorenko et al., 2010) It is important to note that the localizer contrast *sentences* > *nonwords* does not specifically target syntactic processing: instead, it broadly targets high-level language processes, including processing of individual word meanings and combinatorial semantic and syntactic processing (for further discussion, see Fedorenko et al., 2010; Fedorenko et al., 2012b) (see also section 3.4.2).

Condition order was counterbalanced across runs and participants. Experimental blocks lasted 24s (with 5 trials per block), and fixation blocks lasted 16s. Each run (consisting of 3 fixation blocks and 12 experimental blocks) lasted 336s. (Each run contained 4 blocks per condition: in addition to sentences and nonwords, the experiment included a third condition – lists of unconnected words – which was included due to its relevance to another study that was run in the same session.) Each participant completed 3 runs.

3.2.3.2. Critical task

Participants performed a sentence-picture matching task in an event-related design. On each trial, they saw two pictures – each including two characters interacting in some way – and heard a question prompting them to choose one of the pictures (Figure 3.1), by pressing one of two buttons. Sentences contained either syntactically simpler subject-extracted relative clauses (e.g., WHERE IS THE STAR THAT IS GREETING THE CIRCLE?) or syntactically more complex object-extracted relative clauses (e.g., WHERE IS THE CIRCLE THAT THE STAR IS GREETING?).

For the pictures, we used 8 humanized simple shapes (a circle, a square, a triangle, a rectangle, an oval, a heart, a star and an arrow) and 7 easily depict-able actions (chasing, greeting, hugging, lifting, pulling, pushing and touching). Eight characters allowed for 28 unique character-pairs. These were distributed across the 7 actions with 4 character-pairs per action such that (i) each character was used once for each of the 7 actions, and (ii) each action was paired with each of the 8 characters. For each action/character-pair set, we created two versions of a picture (e.g., a circle greeting a star, and a star greeting a circle, as in Figure 3.1), for a total of 56 pictures. The position of the agent relative to the patient (on its left / right) was balanced across items.

For each pair of pictures (e.g., a circle greeting a star, and a star greeting a circle), four sentences were constructed (two per condition), as in (2):

(2) a. Subject-extracted #1: WHERE IS THE CIRCLE THAT IS GREETING THE STAR?
b. Object-extracted #1: WHERE IS THE STAR THAT THE CIRCLE IS GREETING?

c. Subject-extracted #2: WHERE IS THE STAR THAT IS GREETING THE CIRCLE?

d. Object-extracted #2: WHERE IS THE CIRCLE THAT THE STAR IS GREETING?

The sentences were recorded by a female native speaker of English, with a natural prosody, which was created to be as similar as possible across trials and conditions.

Each participant saw each pair of pictures four times over the course of the experiment, twice in the subject-extracted condition, and twice in the object-extracted condition. Pairs of pictures were distributed across four runs such that there was only one occurrence of each pair of pictures per run. So, across the experiment there were a total of 28 picture pairs (as in Figure 3.1) × 4 versions of a sentence = 112 trials (56 trials per condition). The position of the target picture (left, right) was randomized across trials.

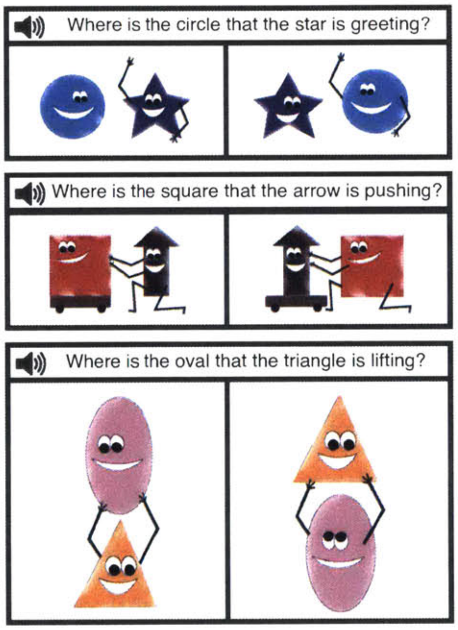

Figure 3.1. Schematic illustration of sample trials in the object-extracted condition. In these instances, the picture matching the sentence is on the left.

Trials were 6s long. Each trial began with a 200ms fixation, followed by the presentation of the pictures and the sentence. The pictures were presented for 4000ms, followed by an extra 1800ms of fixation. Sentence onset was simultaneous with picture onset and each sentence lasted between 4510ms and 5373ms (*M*=4919ms). Participants could respond as soon as the sentence began and through the end of the trial. Each of the four runs lasted 252s, which included 28 6s trials and 84s of fixation (interleaved among the trials, such that the inter-trial interval varied between 0 and 8s). Four condition orders were created using the freely available *optseq* algorithm (Dale, 1999). These orders varied across runs.

3.2.4. Data preprocessing

MRI data were analyzed using SPM5 and custom Matlab scripts (available in the form of an SPM toolbox from http://www.nitrc.org/projects/spm_ss). Each participant's data were motion corrected and then normalized into a common brain space (the Montreal Neurological Institute (MNI) template) and resampled into 2mm isotropic voxels. The data were then smoothed with a 4mm Gaussian filter and high-pass filtered (at 200s). For both the localizer task and the critical task, effects were estimated using a general linear model (GLM) in which each experimental condition was modeled with a boxcar function convolved with the canonical hemodynamic response function. The boxcar function for the localizer task modeled entire blocks; the function for the critical task modeled entire trials.

3.2.5. Data analysis

3.2.5.1. Traditional group analysis

Prior to conducting our key analyses, we aimed to replicate prior findings that used group analyses and reported activations for syntactic complexity manipulations mostly in the inferior frontal gyrus and posterior middle temporal gyrus. Therefore, we ran a group analysis of our critical task by: (i) creating a whole-brain, syntactic complexity contrast map for each participant, contrasting the GLM beta-weights for the object-extracted condition with the weights for the subject-extracted condition; and (ii) entering the contrast maps of all participants into a second-level GLM analysis (p<0.001, uncorrected).

3.2.5.2. Group-constrained, Subject-Specific analysis

Unlike the previous fMRI investigations of syntactic complexity that used traditional group analyses, our key analyses here were performed within regions of interest that were defined functionally in each individual participant. These regions of interest were defined using the *sentences > nonwords* contrast in the language localizer task. To do so, we used the group-constrained, participant-specific analysis method developed in Fedorenko et al. (2010; Julian et al., 2012). In particular, functional regions of interest (fROIs) were constrained to fall within a set of functional "masks" which indicated the expected gross locations of activations for this contrast and which were generated based on a group-level data representation from an independent group of participants (see Figure 3.2A; Fedorenko et al., 2010). These masks were intersected with each individual participant's activation map for the *sentences > nonwords* contrast, and the voxels falling within each mask were sorted based on their *t*-values for the localizer contrast, choosing the top 10% of voxels as that participant's functional region of interest (see Figure 3.2c for sample fROIs). This top *n%* approach ensures that the fROIs can be defined in every participant – thus enabling us to generalize the results to the entire population (Nieto-Castañón and Fedorenko, 2012) – and that fROI sizes are the same across participants. However, qualitatively similar results were obtained in an alternative analysis approach where the fROIs were defined as all the voxels that (i) fell within the relevant mask, and (ii) passed a fixed significance threshold (p<0.001, uncorrected) at the whole-brain level.

Eight fROIs were defined in each participant. These included three fROIs on the lateral surface of the left frontal cortex in the inferior frontal gyrus (IFG) and its orbital part (IFGorb) as well as in the middle frontal gyrus (MFG); and five fROIs on the lateral surface of the temporal and parietal cortex, in the anterior temporal cortex (AntTemp), middle anterior temporal cortex (MidAntTemp), middle posterior temporal cortex (MidPostTemp), posterior temporal cortex (PostTemp) and angular gyrus (AngG). These are "core" regions in the left hemisphere which most robustly and consistently emerge in the investigations of the language network and which include the regions most frequently linked to syntactic processing (but see Appendix 3A for information on the responses to syntactic complexity of the right hemisphere homologues of these regions, and a couple of additional brain regions that consistently emerge for the localizer contrast).

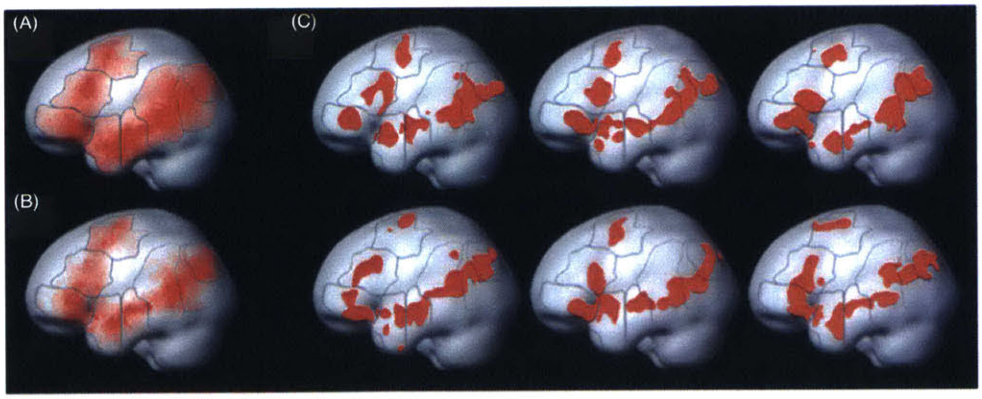

Figure 3.2. Functional regions of interest (fROIs) in the language network. (A) The probabilistic overlap map for the contrast *sentences > nonwords* in a prior dataset of 25 participants (experiments 1 and 2 in Fedorenko et al., 2010). This map was used for generating group-based masks (outlined in gray) that were then used in the current experiment to constrain the selection of individual participants' fROIs. (B) The probabilistic overlap map of individual fROIs in the current experiment (shown in red), constrained to fall within the masks (outlined in gray) that were defined based on the prior data shown in (A). (C) Individual fROIs in six sample participants in the current experiment.

To estimate the responses of fROIs to the conditions of the language localizer, we used an across-runs cross-validation procedure. In particular, each participant's activation map was first computed for the *sentences* > *nonwords* contrast using all but one run of data, and the 10% of voxels with the highest *t*-values within a given mask were selected as that participant's fROI. The response of each fROI to the same contrast was then estimated using the left-out run. This procedure was iterated across all possible partitions of the data, and the responses were finally averaged across the left-out runs to derive a single response magnitude for each condition in a given fROI/participant. This *n*-fold cross-validation procedure (where *n* is the number of functional runs) allows one to use all of the data for defining the fROIs and for estimating their responses (Nieto-Castañón and Fedorenko, 2012), while ensuring the independence of the data used for fROI definition and for response estimation (Kriegeskorte et al., 2009). To estimate the responses of fROIs to the conditions of the critical experiment (i.e., to object-extracted and subject-extracted sentences), data from all runs of the language localizer experiment were used for defining the fROIs.

To summarize the logic of our approach: the language localizer allows us to identify a set of voxels / regions that respond robustly during language processing. We then focus specifically on these regions to test their responses to the critical contrast between syntactically more complex (object-extracted) and syntactically simpler (subject-extracted) relative clauses. If syntactic processing is distributed across the entire language network – rather than localized to particular regions – we should find (i) sensitivity to syntactic complexity in most or all of our fROIs, and (ii) no region-by-condition interactions, indicating that the different regions are similarly sensitive to syntactic complexity.

Statistical tests across participants were performed on the beta weights extracted from the fROIs as defined above. Two contrasts were examined: (i) *sentences* > *nonwords*, and (ii) *object-extracted* > *subject-extracted.*

3.3. Results

3.3.1. Behavioral results

Due to a script error, behavioral responses for the sentence-picture matching task were not recorded for 6 of the 13 participants. However, 4 of these participants were later brought in and re-tested behaviorally on exactly the same version of the experiment, so that altogether we obtained behavioral data from 11 of the 13 participants. In those 11 participants, we replicate the standard complexity difference, with slower RTs (4.57s vs. 4.40s; $t_{(10)}=2.23$, $p<0.05$) and lower accuracies (91% vs. 96%; $t_{(10)}=-2.95$, $p<0.05$) in the object-extracted condition (see also section 3.3.3.6 below).

3.3.2. Traditional fMRI group analysis

A whole-brain, random-effects group analysis of the syntactic complexity contrast *object-extraction > subject-extraction* revealed several significant activation clusters, of which three appeared in regions of the left-hemisphere commonly associated with language processing. These clusters were located in (i) the posterior part of the middle temporal gyrus; (ii) the triangular part of the inferior frontal gyrus; and (iii) the mid-anterior part of the superior temporal gyrus (Table 1). The former two clusters are broadly consistent with the activation foci most commonly found in prior studies using similar contrasts. This can be seen in Figure 3.3, showing our group-level activation map along with marked locations of previously reported syntactic complexity effects.

Table 3.1. Activation clusters for the syntactic complexity contrast (*object-extracted > subject-extracted*) identified with traditional group analysis

	Center coordinates (mm)[b]			**Volume**	**Peak**
Region[a]	**x**	**y**	**z**	**(mm^3)**	***t*-value**
Left middle temporal gyrus (posterior)	-48	-58	5	4624	7.09*
Left inferior frontal gyrus, triangular	-40	28	0	88	5.36*
Left superior temporal gyrus (mid-anterior)	-57	-12	-5	280	7.26*

[a] We only report clusters that are located in left-hemispheric regions commonly associated with the language network.
[b] Coordinates are reported in MNI space.
* $p<0.001$, uncorrected for whole-brain multiple comparisons.

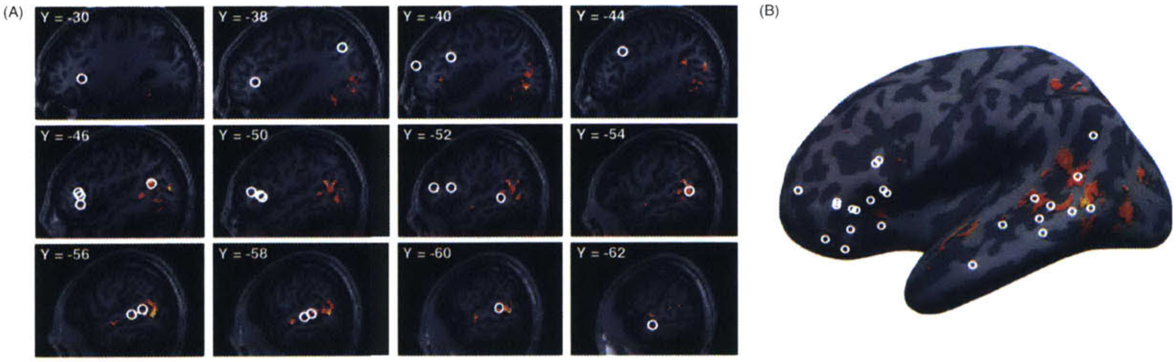

Figure 3.3. Syntactic complexity effects in the left hemisphere identified with traditional group analysis. Both (A) and (B) show the activation map of our critical contrast, *object-extraction > subject-extraction* ($p<0.001$, uncorrected for whole-brain multiple comparisons) in hot colors. White circles show the locations of activations to similar syntactic complexity contrasts as reviewed by Friederici (2011; referred to in that paper as "studies of movement"). Activations in the current study fall within the same general locations found previously, namely, the posterior middle temporal gyrus and the inferior frontal gyrus. (A) The effects are superimposed on sagittal slices of an anatomical scan from one of our participants. (B) The effects are projected onto an inflated cortical surface of an average brain in MNI space.

3.3.3. Group-constrained, participant-specific fMRI analysis

3.3.3.1. Are syntactic complexity effects localized to particular regions within the language network?

Replicating previous work, we find robust responses for the localizer contrast (*sentences* > *nonwords*) in each of the eight fROIs, using across-run cross-validation (for all regions, $t_{(12)}>5$, $p<10^{-4}$; *t*-tests are across participants). Critically, all of the regions defined by the localizer, except for the AngG fROI, also showed a significant effect for our syntactic complexity manipulation. All effects remain significant after false-discovery rate (FDR) correction for the number of fROIs. These key results are summarized in Figure 3.4.

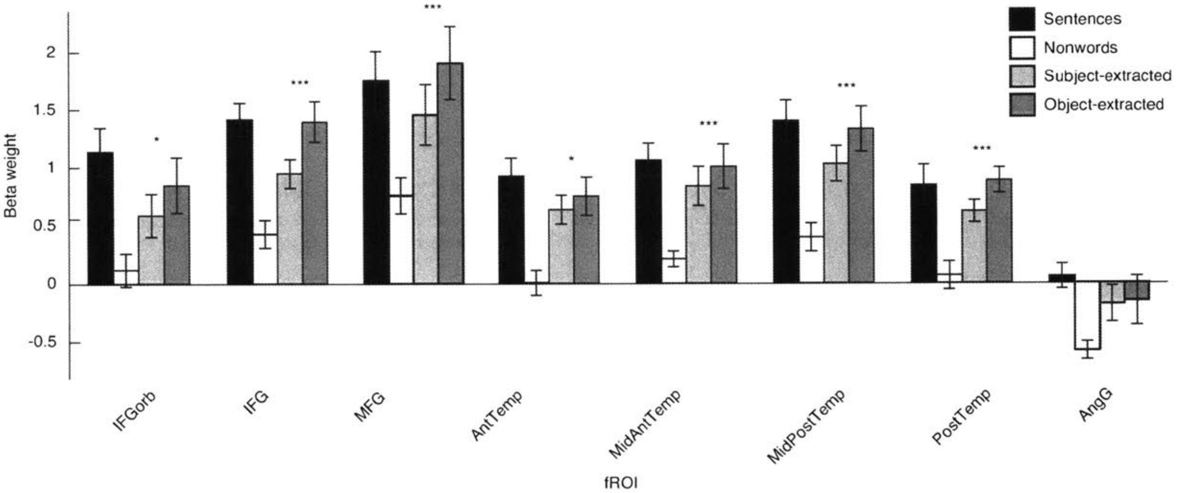

Figure 3.4. Responses of the language fROIs to the conditions of the language localizer and the critical experiment. Error bars represent standard errors of the mean by participants. The *sentences* > *nonwords* contrast is highly significant ($p<10^{-4}$) in every region (this analysis was carried out using across-runs cross-validation, so that the data used to define the fROIs and estimate the responses are independent, as described in section 3.2.5.2). For the *object-extracted* > *subject-extracted* contrast: * significance at the $p<0.05$ level, and *** significance at the $p<10^{-3}$ level or stronger. All effects remain significant after an FDR correction for the number of regions ($n=8$). (Note that it is difficult to directly compare the magnitudes of response to the sentences condition of the localizer task and the magnitudes of response to the two critical conditions, because of many differences in the design, materials and procedure across the two experiments.)

3.3.3.2. Do regions differ reliably with respect to how sensitive they are to syntactic complexity?

We found that every region within the language network (with the exception of the AngG fROI) responds reliably more strongly during the syntactically more complex object-extracted condition than during the syntactically simpler subject-extracted condition. However, the difference between these two conditions is numerically larger in some regions than others. In particular, the largest syntactic complexity effects are observed in the brain regions that have been reported most consistently in previous studies (i.e., regions in and around Broca's area and regions in the posterior temporal cortex; the MFG fROI – also reported in a few prior studies (e.g., Meltzer et al., 2010) – also shows a large effect). One possibility is that these regions are in fact more strongly engaged in – and thus perhaps play a bigger role in – syntactic processing compared to the rest of the language network. Nevertheless, our data suggest that this is not the case.

In particular, the overall response to language (e.g., the response to the sentences condition of the language localizer relative to the fixation baseline) also varies across regions: it is numerically stronger in the more superior and dorsal frontal regions (the IFG and MFG fROIs) than in the inferior and ventral IFGorb fROI, and it is stronger in the MidPostTemp fROI than in the more anterior temporal regions (AntTemp and MidAntTemp fROIs) and the more posterior temporal/temporo-parietal regions (PostTemp and AngG fROIs). This pattern of different-strength BOLD responses across the language regions is consistent across participants and studies (Fedorenko et al., 2010; Fedorenko et al., 2011; Mahowald and Fedorenko, 2016)

Given that effect sizes tend to scale with overall response strength, it is not surprising that the effects of syntactic complexity are more difficult to detect in brain regions where the response to language is overall weaker. Indeed, we find that the overall language response (*sentences > fixation* effect) in a particular fROI is a significant predictor of that fROI's response to syntactic complexity (*object-extracted > subject-extracted* effect), using a linear mixed-effects regression predicting the syntactic effect size from the overall response with random intercepts and slopes for fROI and participant (β=0.18, t=3.16, $\chi^2_{(1)}$=7.39, p<0.01). Note that this finding cannot be accounted for by differences in the number of voxels across fROIs, because we obtained similar results when equating the volumes of our fROIs (β=0.18, t=3.43, $\chi^2_{(1)}$=8.46, p<0.01).

In Figure 3.5, we show the relationship, across the eight ROIs, between the overall response to language (*sentences > fixation*, in the localizer experiment) and the size of the *object-extracted > subject-extracted* effect (effect sizes are averaged across participants). As can be seen, the PostTemp and IFG fROIs both show a larger syntactic effect than would be predicted by their overall language response (they fall above the trend line), whereas AntTemp and MidAntTemp fROIs show a smaller syntactic effect than would be predicted by their overall language response. However, none of these deviations are significant. Specifically, allowing the association between overall language response and syntactic complexity effect size to vary across fROIs, by including a random "overall language response" slope for each fROI, does not significantly improve the model ($\chi^2_{(2)}$=0.25, n.s.). The standard deviation of this random "overall language response" slope is very small (0.01) compared to the size of the corresponding fixed effect (β=0.18), suggesting that our different fROIs contribute indistinguishable data to the model. Thus, although language regions may differ slightly in the relative strengths of the syntactic complexity effect, most of this variance appears to be accounted for by differences in the overall response to language stimuli across the language network. Beyond this explainable variance we find no evidence for a region-by-condition interaction and we cannot reject the hypothesis that our fROIs are all similarly sensitive to syntactic complexity manipulations.

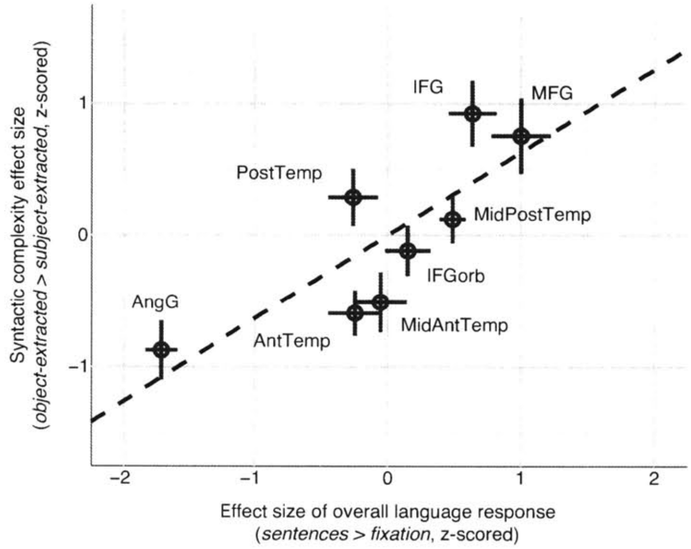

Figure 3.5. The syntactic complexity effect size co-varies with overall sensitivity to language. The mean size, across participants, of the syntactic complexity effect (*object-extracted* > *subject-extracted*) is plotted against the mean effect size of the overall response to language, as estimated in the localizer experiment (*sentences* > *fixation*). To control for inter-individual differences in the overall response strength, data for the eight fROIs were *z*-scored within each participant prior to averaging. Crosses show standard errors across participants for both effects. A dashed, black line depicts the linear regression line for predicting the syntactic complexity effect based on the overall language response, and was estimated for visualization purposes only (the linear mixed-effects regression reported in section 3.3.3.2. was carried out using individual data from all participants).

3.3.3.3. The spatial pattern of syntactic complexity effects is better explained by language-specific responsiveness than by general, non-specific proneness to signal loss

Our analysis of the relationship between the size of the syntactic complexity effects and overall response magnitude to language was performed on contrast estimates that were averaged across voxels in each fROI. This averaging might have obscured potential heterogeneity within these regions. It is therefore possible that, on a finer-grain spatial scale, one would not find associations between syntactic complexity effect sizes and overall language response. To test this possibility, we here explore the relationship between the two effects across individual voxels.

A correlation between the syntactic complexity effect size and overall language response across voxels would be compatible with two interpretations. One possibility is that the association is not language-specific: a strong correlation across voxels would be expected for any two effects, linguistic or non-linguistic, due to physiological artifacts. In particular, inter-regional differences in vascularization (Harrison et al., 2002; Ances et al., 2008; Ekstrom, 2010; Wilson, 2014) or proneness to signal loss (Jezzard and Clare, 1999; Menon and Kim, 1999) might explain why different contrasts co-vary across voxels (e.g., regardless of the particular contrast, effect sizes across voxels might scale with the voxels' distance from air-tissue interfaces). An alternative interpretation, however, is that the association between the two effects is language-specific and would not generalize to non-linguistic effects.

To distinguish these possibilities, we ran a model predicting the size of the syntactic complexity effect (*object-extracted* > *subject-extracted*) across individual voxels using two predictors: a non-linguistic effect and a language-specific response (*sentences* > *fixation* in the localizer task). Our non-linguistic effect contrasted two versions of a spatial working memory task differing in difficulty (*hard* > *easy*). In this task, which our participants performed in the scanner for another study, participants have to keep track of four vs. eight locations in a 3×4 grid. This task has previously been shown to have reliable variability across cortical voxels (allowing, in particular, the functional localization of frontal and parietal regions of the "cognitive control" or "multiple demand" network; Fedorenko et al., 2013b). According to the first interpretation above, predicting the size of the syntactic complexity effect from the size of the non-linguistic effect would not benefit from adding the language-specific response magnitude as a predictor (given that all contrasts should show strong correlations). However, according to the second interpretation, the size of the syntactic complexity effect would be predicted by the size of the language-specific response magnitude above and beyond the non-linguistic effect size.

A linear, mixed-effects regression model with random intercepts and slopes for both participant and fROI supported the second interpretation: the contribution of the language-specific response magnitude to the model was significant (β=0.16, t=5.22, $\chi^2_{(1)}$=15.7, $p<10^{-4}$). In fact, when the language-specific response magnitude was included in the model, the non-linguistic effect size was not a significant predictor of the syntactic complexity effect size (β=-0.03, t=-0.31, $\chi^2_{(1)}$=0.09, n.s.). We note that the localizer runs used to define fROIs for this model were the same runs in which the language-specific response magnitude was evaluated; however, there is no non-independence involved in this procedure, because we are evaluating correlations across voxels instead of effect sizes averaged over the chosen voxels.

This lack of association between syntactic complexity and non-linguistic effects is not due to the restricted range of contrast values in our fROIs (where non-linguistic effects are very weak), as our results extend beyond those regions. Specifically, similar results were obtained when we ran the model on all individual voxels falling within the group-based masks instead of only including voxels falling within participant-specific fROIs (contribution of language-specific response magnitude: $\beta=0.15$, $t=5.21$, $\chi^2_{(1)}=17$, $p<10^{-4}$; contribution of non-linguistic effect size: $\beta=-0.01$, $t=-0.11$, $\chi^2_{(1)}=0.01$, n.s.). Therefore, the correlation between the size of the syntactic complexity effect and the response magnitude to language is functionally specific, and does not generalize to non-linguistic contrasts. It is therefore unlikely that physiological artifacts, such as regional differences in vascularization, are the main factor underlying this correlation.

3.3.3.4. Language fROIs are sensitive not only to sentence-level syntax, but also to lexical information

We would like to stress that the above results should not be taken as indication that our localizer contrast (*sentences > nonwords*) is in fact just a localizer for syntactic processing. Our localizer targets regions involved in various aspects of high-level linguistic processing, including both semantic and syntactic processing at both the lexical and sentence levels, as previous work from our lab has shown (Fedorenko et al., 2010; Fedorenko et al., 2012b). To more directly ensure that our localizer did not exclusively target sentence-level syntactic processing, we took advantage of our localizer design which included, besides sentences and nonwords, a third condition: word lists that did not form sentences (and that were included here for the purposes of another study). Below, we briefly report three analyses targeting the contrasts between word-lists and control conditions (either nonword lists or fixation) as a measure of lexical processing unrelated to sentence-level syntax.

First, we measured the size of the *words > nonwords* effect in our fROIs (localized with the *sentences > nonwords* contrast). Replicating our prior work (Fedorenko et al., 2010) and using across-runs cross-validation as in the other analyses, we observed reliable responses in all fROIs except for the MFG (for all tests, $t>1.92$, $p<0.05$; all regions except for the IFG remained significant after FDR correction for the number of regions). Second, we found that across individual voxels in all eight fROIs, the size of the *sentences > fixation* effect was predicted by the size of the *words > fixation* effect, above and beyond the prediction provided by the non-linguistic working memory effect (linear, mixed-effects regression with random intercepts and slopes for both participant and fROI: $\beta=1.18$, $t=16.23$, $\chi^2_{(1)}=36.21$, $p<10^{-8}$). These findings demonstrate that our fROIs are sensitive not only to sentence-level syntactic information, but also to lexical information.

Third, we repeated our main analysis of the syntactic complexity effect (reported in section 3.3.3.1), but now defined fROIs using the *words > nonwords* contrast (instead of using the *sentences > nonwords* contrast). Despite the fact that this alternative localizer tends to produce weaker contrast effects compared to the localizer reported in the paper, we found a significant syntactic complexity effect in all regions except for the AntTemp and AngG (for all tests, *t*>2.83, *p*<0.03; all regions remained significant after FDR correction for the number of regions) (Figure 3.6).

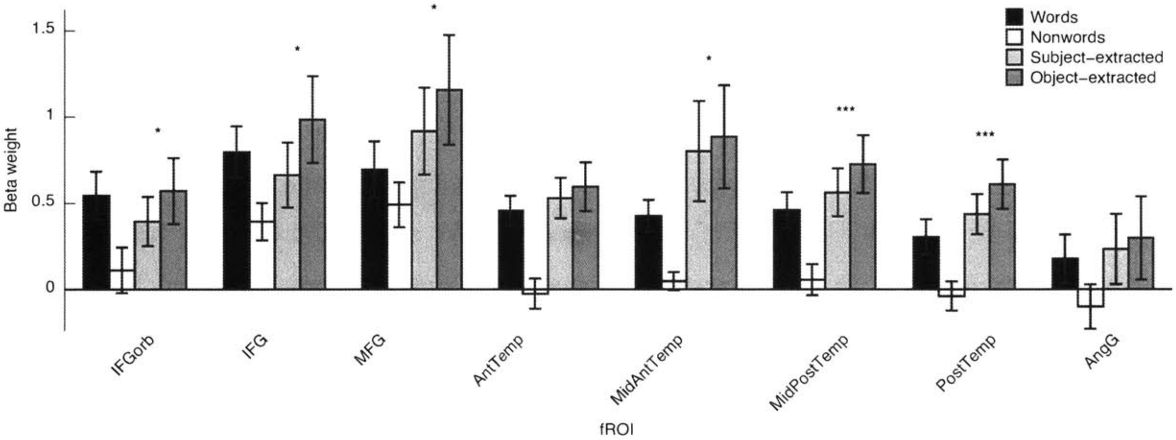

Figure 3.6. Replication of the main result using an alternative language localizer contrast (*words > nonwords*). Responses of the language fROIs are shown to the conditions of the alternative language localizer and the critical experiment (*object-extracted > subject-extracted*), using the same conventions as in Figure 3.4.

The apparent lack of a syntactic complexity effect in the AntTemp (cf. a significant effect reported in the main text, when this fROI was localized with the *sentences > nonwords* contrast) should be interpreted with care. First, the effect sizes of the *words > nonwords* localizer contrast are much weaker than those of the *sentences > nonwords* localizer contrast, throughout the entire language network. Thus, the *words > nonwords* localizer contrast has inferior localization capacities, and this weakness might account for the lack of a syntactic complexity effect in the AntTemp. Second, we emphasize that a functional difference between the AntTemp and the other language regions requires a region-by-condition interaction.

Alternatively, the lack of a syntactic complexity effect in a region localized with a lexical contrast might imply that the anterior temporal lobe contains a sub-region that is recruited for processing word-level information but not sentence-level syntactic information. Such a sub-region is perhaps separate from the sub-region that did show a syntactic complexity effect in the main analysis (for a similar suggestion, see Pascual et al., 2015). This conjecture might explain why the role of anterior temporal lobe in language processing remains debated: on the one hand, it has been reported to engage in syntactic and semantic combinatorial processes above the word level, and only for linguistic stimuli but not for other meaningful stimuli (Humphries et al., 2001;

Vandenberghe et al., 2002; Humphries et al., 2005; Rogalsky and Hickok, 2009; Baron and Osherson, 2011; Bemis and Pylkkänen, 2011; Brennan et al., 2012); on the other hand, it has been identified as an amodal, non-linguistic semantic hub for simple concepts (Patterson et al., 2007; Lambon-Ralph et al., 2010; Wong and Gallate, 2012; Jefferies, 2013; Mesulam et al., 2015).

3.3.5. Language fROIs are sensitive to syntactic manipulations during a passive reading task

Straightforward interpretations of our main result assume that the syntactic complexity effects we observed reflect, in some way or another, an inherent difference between object-extracted and subject-extracted relative clauses – a "pure" difference between construction types that would replicate whenever such sentences are processed. Is it possible that these effects instead result from an interaction between construction and our particular sentence-picture matching task? Perhaps some extra-linguistic aspects of this task are more difficult when hearing object-extracted sentences compared to subject-extracted sentences, accounting for our results.

For instance, it has been previously argued that both sentence types tend to be initially parsed by assigning the active role of an agent to the first noun encountered (a "subject-first" assumption; see Frazier, 1987; Frazier and d'Arcais, 1989; Schriefers et al., 1995; Schlesewsky et al., 2000; Traxler et al., 2002); this assignment is correct in subject-extracted sentences (THE CIRCLE THAT IS GREETING THE STAR), but requires reanalysis in the case of object-extracted sentences, where the first noun is the patient of an action (THE CIRCLE THAT THE STAR IT GREETING). Perhaps then, upon hearing the first noun in our sentence stimuli, participants searched for a picture in which that noun was depicted as the agent rather than the patient. Such a strategy would correctly solve the task for the subject-extracted sentences, but would force participants to switch pictures upon reanalysis of the object-extracted sentences. Some cognitive process involved in this picture switching might underlie the stronger activations in language regions observed for the latter sentences compared to the former.

Interpreting our results as reflecting stimulus-task interactions appears to require that the extra-linguistic differences in task performance for the two sentence types involve executive functions (guiding behavioral strategies), response inhibition, working memory or other, similar, domain-general cognitive resources. However, previous data show that such mental processes do not recruit the language network (see Discussion). Specifically, language regions respond at or below a low-level baseline to tasks that have general demands similar to those of the sentence-picture matching task (Fedorenko et al., 2011; Fedorenko et al., 2012a).

Furthermore, effects of syntactic complexity like the one studied here are among the most robust sentence-level linguistic phenomena and have been shown to hold across a wide range of paradigms in the prior literature (see Materials and methods; e.g., reading with comprehension questions or plausibility judgments, listening with comprehension questions, listening with a concurrent lexical-decision task or nonword detection task, sentence repetition, etc.). It is generally assumed that the mental processes underlying syntactic complexity effects across all these diverse paradigms are the same.

Importantly, syntactic complexity effects also replicate in naturalistic materials under passive reading conditions, where no interaction with an externally imposed task is expected (e.g., Demberg and Keller, 2008) (e.g., Demberg & Keller, 2008). To further support this claim, we wanted to demonstrate that the language fROIs localized in the current study exhibit sensitivity to syntactic processing during a passive reading task. For this purpose, we analyzed data reported in (Fedorenko et al., 2010) from 12 participants who passively read linguistic stimuli, including critically, Jabberwocky sentences (which preserve the word order, function words and functional morphology of real sentences but use nonwords) and lists of unconnected nonwords. Because only the stimuli in the former condition contain identifiable syntactic structure (e.g., due to the presence of function words), we interpret the *Jabberwocky > nonwords* contrast as a syntactic contrast. These data were analyzed with the same procedures described above, and are presented in Figure 3.7. Consistent with our main results, all language fROIs show a stronger response to Jabberwocky sentences than to nonwords (including the AngG, which shows the smallest effect size and a negative beta weight for nonwords).

Figure 3.7. Responses of the language fROIs to real sentences, Jabberwocky sentences and lists of nonwords. The *sentences > nonwords* contrast was used to localize the fROIs. The effect sizes for all three conditions, as well as the *Jabberwocky > nonwords* contrast, were then evaluated in independent data, using across-runs cross-validation. Asterisks denote the significance of the *Jabberwocky > nonwords* contrast. Conventions are the same as in Figure 3.4.

We note that our critical contrast reported in the main text (*object-extracted > subject-extracted*) still provides stronger evidence for sensitivity to syntactic processing compared with the *Jabberwocky* > *nonwords* contrast, because the former (i) is "tighter" and based on a minimal pair of sentences that contain identical words and only differ in their word order (syntax); and (ii) uses sentences with real words, as opposed to the less natural nonwords and, therefore, has stronger ecological validity. Nevertheless, the *Jabberwocky* > *nonwords* contrast reported here contributes converging evidence in support of our main result. We believe that these data, measured during a passive-reading task, alleviate the concerns about our results reflecting an interaction between sentence-type and the sentence-picture matching task. We note that a subject-first strategy could still underlie our effects, if such an account rested on the linguistic consequences of this assumption (consistent with a frequency-based interpretation of our results, advocated in section 3.4.1.2) instead of its extra-linguistic and task-specific consequences.

3.3.6. Does the size of the fMRI syntactic complexity effect predict any aspects of behavioral performance?

Participants varied with respect to the size of the *object-extracted* > *subject-extracted* effect. Might these differences be functionally important? In an exploratory analysis, we examined the relationship between the size of the syntactic complexity effect and behavioral performance. Before doing so, we examined the reliability of the behavioral measures, because if the performance estimates are too noisy at the individual-participant level, there is no reason to expect them to correlate with the effect size in fMRI data. We considered four measures: difference in the accuracy / RT between the object-extracted and subject-extracted conditions, and overall (averaged across the two conditions) accuracy / RT. For each measure, we split the data into odd and even-numbered trials and looked for correlations across participants. Although both of the RT measures were highly reliable (r=0.91 for the overall RT, and r=0.51 for the object-extracted vs. subject-extracted difference), neither showed a reliable relationship with the size of the fMRI syntactic complexity effect. As for accuracies, the difference measure was not correlated between odd- and even-numbered trials (r=-0.14), but the overall accuracy was highly reliable across the two data halves (r=0.70).

When we correlated the overall behavioral accuracies with the size of the fMRI syntactic complexity effect, we found that participants with larger syntactic complexity effects in the fMRI BOLD signal performed significantly worse in the task. As we see in Figure 3.8, all 8 fROIs show this trend, again highlighting the similarity among these regions with respect to their engagement in syntactic processing. The effect is significant in a linear, mixed-effects model predicting the size of the fMRI syntactic complexity effect from the logit-transformed accuracies with random intercepts for participant and random intercepts and slopes for fROI (β=-0.21,

t=-3.48, $\chi^2_{(1)}$=9.29, p<0.01). This relationship remains significant after a Bonferroni correction for the number of behavioral measures examined (n=4).

One way to interpret this relationship is in terms of comprehension efficiency. For example, some participants may have greater exposure to syntactically complex object-extracted structures and/or have greater working memory capacity (see section 3.4.1.2), and consequently may not need to activate their language regions more strongly to process the more syntactically complex structures. Such participants are also likely to be overall better in their language comprehension ability, thus answering comprehension questions more accurately.

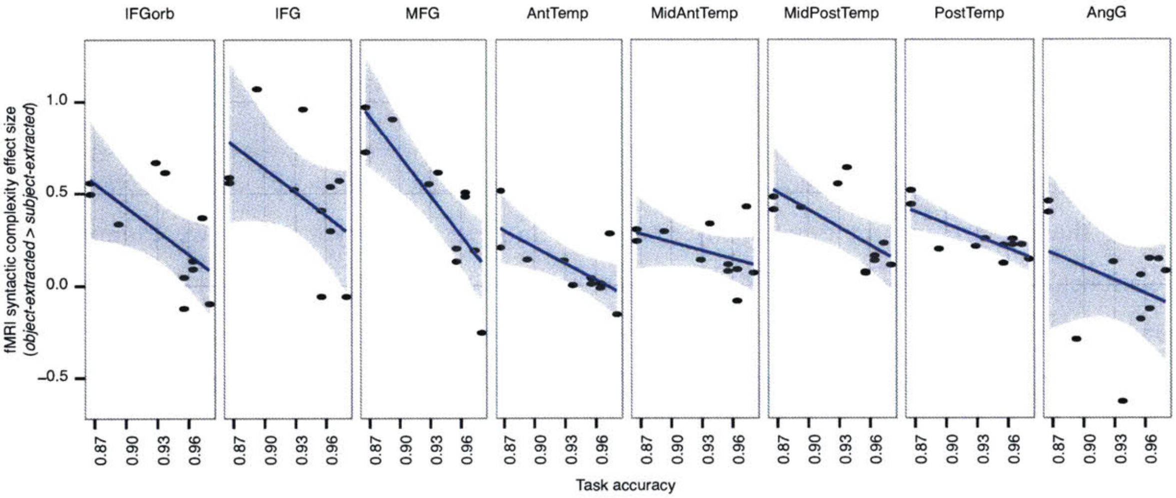

Figure 3.8. The relationship between task accuracy and the size of the syntactic complexity effect (*object-extracted > subject-extracted*) in fMRI. Data is shown for each of the 8 fROIs, which all show a downward trend. Blue lines are based on a simple linear regression for each region, with smoothed 95% confidence intervals shaded in gray. Most of the points fall above 0, which shows the main effect of increased fMRI response to the object-extracted condition relative to the subject-extracted condition.

3.3.4. Why do the traditional group-analysis and the group-constrained, participant-specific analysis produce different results?

Traditional group analyses, by design, identify regions of activation that are overlapping across many participants. Regions in which activations show less inter-individual overlap will therefore be missed by such analyses (see, e.g., Nieto-Castañón and Fedorenko, 2012 for underlying mathematics and simulation data). Our alternative analysis method, in contrast, allows for some variability in the locations of activations across people due to its use of individually defined functional regions of interest. Given our results above, we reasoned that activation maps for the syntactic complexity effect would show relatively higher inter-individual overlap in the IFG and MidPostTemp compared to other fROIs, as these regions were identified by the group analysis.

To visualize these potential differences across regions, we identified regions showing syntactic complexity effects in each participant and evaluated their inter-individual overlap. To this end, we first identified activation peaks in individual maps of the syntactic complexity contrast (*object-extracted* > *subject-extracted*) using a watershed algorithm (to prevent the algorithm from over-parcellating the contrast maps, they were smoothed with a 8mm FWHM Gaussian kernel). An "activation neighborhood" around each peak was then defined (in the original, non-smoothed activation maps) as the largest contiguous set of surrounding voxels having numerically positive contrast estimates. Finally, for each voxel in common MNI space, we counted the number of participants for whom that voxel belonged to an activation neighborhood.

Figure 3.9 shows the overlap measures we obtained. As can be seen, activations in IFG and MidPostTemp show the highest overlap across participants, along with MFG. These three regions are also the ones where the syntactic complexity effects are numerically the strongest, and these two observations are plausibly linked. Given that regions of the language network show

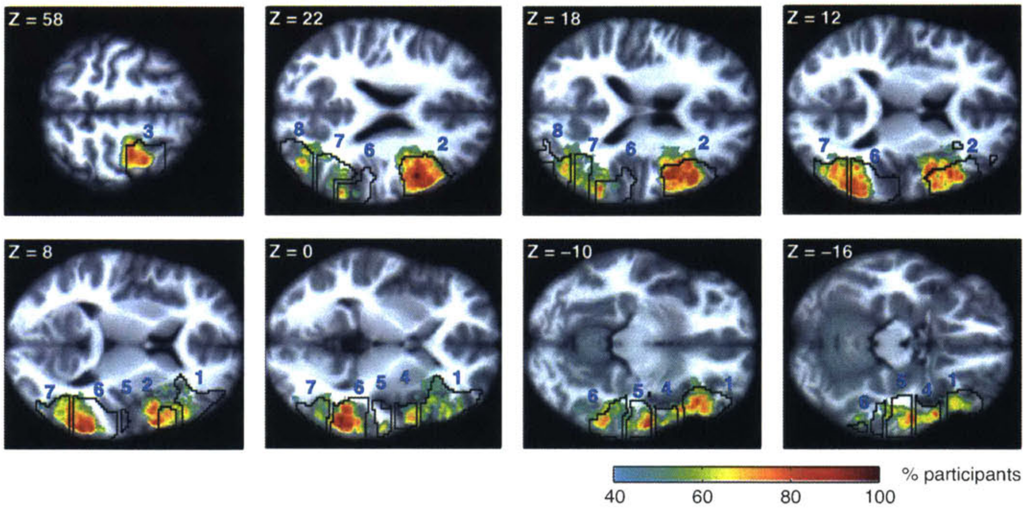

Figure 3.9. Overlap across participants in the anatomical location of the syntactic complexity effect. Heat maps depict voxels in which more than 40% of participants have an "activation neighborhood" for the syntactic complexity contrast (*object-extracted* > *subject-extracted*). Neighborhoods were defined as maximal sets of contiguous voxels that surrounded an activation peak and had contrast estimates numerically greater than zero. Black contours depict the group-based masks (Fedorenko et al., 2010) used to define fROIs. fROI numbering follows Figure 3.4: 1, IFGorb; 2, IFG; 3, MFG; 4, AntTemp; 5, MidAntTemp; 6, MidPostTemp; 7, PostTemp; 8, AngG. Data are superimposed on horizontal slices of an average T1 scan in common MNI space. Slices were chosen to maximize visibility of the greatest overlap in each mask. Note the especially high overlap (dark red color) in the IFG and MidPostTemp.

relatively high inter-individual variability in their functional-to-anatomical mapping (Amunts et al., 1999; Fischl et al., 2008; Frost and Goebel, 2012; Tahmasebi et al., 2012), overlap in activation maps across participants is mainly expected in the most responsive regions that have high peaks and, thus, larger activation neighborhoods.

3.4. Discussion

Our results demonstrate that syntactic complexity effects – greater responses to more syntactically complex sentences – are not localized to particular regions within the language network, but are instead found throughout the entire network. Although our results are consistent with prior studies that have observed these effects in the inferior frontal and posterior temporal brain regions, we also show that these effects obtain in the rest of the language regions (with the sole exception of the AngG fROI), including the language-responsive regions in the orbital IFG, and in the anterior portions of the lateral temporal cortex.

As discussed in the Introduction, an architecture where syntactic resources are distributed across the language network fits well with the findings from the patient literature: deficits in syntactic comprehension have been reported following damage to many different components of the language network (Caplan et al., 1996; Dick et al., 2001; Wilson and Saygin, 2004; Mesulam et al., 2015) (see Mesulam, 1990, for an early discussion of distributed language processing).

Some previous neuroimaging results further support the idea of distributed syntactic processing, although they do not isolate syntactic processing from other aspects of language comprehension, as the current manipulation does. For example, a contrast between sentences (in which words combine to form syntactic structures) and lists of unconnected words (devoid of such structures) produces activation across the language network (Snijders et al., 2009; Fedorenko et al., 2010; Bedny et al., 2011; Pallier et al., 2011; Brennan and Pylkkänen, 2012) (see also earlier studies: Mazoyer et al., 1993; Schriefers et al., 1995; Vandenberghe et al., 2002; Humphries et al., 2005; Humphries et al., 2006; although those typically found activations only in parts of the language network). Admittedly, the *sentences > word lists* contrast is not a "pure" syntactic manipulation, because sentences differ from word lists in additional ways: they also engage compositional semantic processes and possibly, at least for auditory presentation, prosodic processes. A somewhat syntactically purer contrast, between Jabberwocky sentences and lists of unconnected nonwords, also produces a response throughout the language network (Fedorenko et al., 2010; Bedny et al., 2011) (see also section 3.3.3.5).

Nonetheless, a vast majority of prior neuroimaging studies of syntactic complexity have instead supported a localized, rather than distributed, view of syntactic processing. We have argued that these prior investigations may have not observed effects in some parts of the language network because of the poor sensitivity of traditional group-based analyses (Nieto-Castañón and Fedorenko, 2012) and because those regions have an overall weaker response to language and thus smaller, harder to detect, effects, especially for subtle manipulations. Our current findings support this claim: first, we directly contrasted a traditional group-based analysis that found evidence for a few localized foci of syntactic complexity effects, with an analysis based on individual localization of language-responsive fROIs that found these effects robustly present throughout the language network. Second, the regions that the group-based analysis failed to identify appeared to have higher inter-individual variability (i.e., less overlap) in activations. This poorer overlap appeared to coincide with lower responsiveness to language in those regions, compared to the regions that the group-based analysis did successfully identify.

Similar reasoning applies to other studies that have targeted syntactic processing and reported effects only in the inferior frontal and posterior temporal regions (e.g., syntactic violation manipulations: Embick et al., 2000; Cooke et al., 2006; Friederici et al., 2010; Herrmann et al., 2012; or syntactic priming: Santi and Grodzinsky, 2010; Menenti et al., 2012; Segaert et al., 2012). We hypothesize that those effects (and possibly other, non-syntactic, effects; e.g., Devlin et al., 2000), like the syntactic complexity effects studied here, are actually present throughout the language network.

It is worth noting that, contra proposals about Broca's area and parts of the posterior temporal cortex being the core syntactic centers of the brain, a number of researchers have argued that parts of the anterior lateral temporal cortex are instead critically engaged in combinatorial syntactic (and/or semantic) processing (Humphries et al., 2001; Vandenberghe et al., 2002; Humphries et al., 2005; Rogalsky and Hickok, 2009; Baron and Osherson, 2011; Bemis and Pylkkänen, 2011; Brennan et al., 2012; Zhang and Pylkkanen, 2015) (see also section 3.3.3.4). We suspect that, as with the above studies, the observed effects are present across the language network, although it is not at present clear why these studies differ from the studies above in observing the effects in the anterior temporal as opposed to inferior frontal and posterior temporal regions.

3.4.1. Syntactic complexity effects: interpretations and limitations

3.4.1.1 Causal involvement in syntactic processing

The finding that a distributed set of language regions are all sensitive to syntactic complexity manipulations should not be interpreted as demonstrating that all of these regions play an equal role in syntactic processing. For example, we do not suggest that every region that shows a stronger response to syntactically complex sentences than to syntactically simpler sentences is causally involved in syntactic processing. A possible alternative is that only a subset of our language fROIs are critical for processing syntax, but their output rapidly travels to the rest of the language network and is therefore reflected in the temporally slow BOLD signal. Furthermore, whereas syntactic complexity primarily modulates syntactic processing, it may additionally modulate other comprehension processes (like those related to the processing of information structure; e.g., Jackendoff, 1972), thus leading in some fROIs to linguistic but non-syntactic "secondary effects" masking as syntactic effects (but see section 3.4.2).

Like all fMRI studies, our current study is not designed to (and could not) distinguish regions that are causally involved in syntactic processing from those that are more epiphenomenally recruited. Neuropsychological studies are also limited in their ability to identify such distinctions among regions, because naturally occurring brain damage typically encompasses multiple functional areas, as well as extending to white matter tracts connecting regions that may themselves be unaffected by the lesion (Mesulam et al., 2015). Identifying regions that are critical for syntactic processing ultimately requires causal measurements with both high temporal and spatial resolution, such as invasive stimulation studies using subdural electrodes (inserted pre-surgically for medical reasons; e.g., Ojemann et al., 1989). Nonetheless, we emphasize that our current contribution is the demonstration that signals reflecting the modulation of neural activity by syntactic complexity (whatever such activity reflects at the mechanistic and cognitive levels) are present throughout the language network, contrary to many previous suggestions.

3.4.1.2. Experience-based versus working-memory-based accounts of syntactic complexity

For brain regions that are causally linked to syntactic processing – whatever subset of the language network these may turn out to correspond to – another question arises: which of the factors underlying the complexity difference between object-extracted and subject-extracted structures modulate the activity of these regions? Two classes of proposals have been advanced to account for such complexity differences: experience-based theories (Hale, 2001; Gennari and MacDonald, 2008; Levy, 2008; Gennari and MacDonald, 2009; Wells et al., 2009) and working-memory-based theories (for overviews, see O'Grady, 2011; Gibson et al., 2013). According to the former, object-extractions are more difficult to understand because they are less frequent in the input. According to the latter, processing object-extractions places greater demands on working memory because one of the dependents of the verb has to be retrieved from memory when the verb is encountered (Gibson, 1998; Gordon et al., 2001; McElree et al., 2003; Gordon et al., 2004; Grodner and Gibson, 2005; Fedorenko et al., 2006; Lewis et al., 2006; Fedorenko et al., 2013a). Neither class of proposals can fully explain the rich empirical picture that has emerged from dozens of sentence processing studies, and most researchers now agree that a complete account of language comprehension requires both a probabilistic grammar component and a (plausibly domain-general) working memory resource (Lewis et al., 2006; Demberg and Keller, 2008; Boston et al., 2011; Gibson et al., 2013; Levy et al., 2013).

Does the syntactic complexity effect we observed throughout the language network reflect the differences in frequency between object-extracted and subject-extracted constructions, or the different demands they place on working memory? Fedorenko et al. (2011) showed that regions of the language network do not respond to general working memory demands, although in the left frontal cortex they lie adjacent to other, distinct regions that are strongly modulated by working memory demands (Fedorenko et al., 2012a, 2013b). The effects reported here in the language regions are thus unlikely to reflect differences in working memory (cf. Fiebach et al., 2001; Kaan and Swaab, 2002; De Vries et al., 2008; Rogalsky and Hickok, 2011). We therefore conjecture that these effects reflect differences in the relative frequencies of the two constructions, although we note that the design of the current experiment cannot provide evidence favoring either account over the other.

This interpretation of our results does not contradict the contribution of general working memory resources to the syntactic complexity effects. Namely, whereas we have here focused on the regions of the language network, syntactic complexity manipulations also produce responses in the regions of the domain-general fronto-parietal "multiple demand (MD)" network (e.g., Barde et al., 2012), and damage to some MD regions can lead to difficulties with syntactically complex structures (e.g., Amici et al., 2007). More generally, MD regions respond to diverse executive tasks (Duncan and Owen, 2000; Corbetta and Shulman, 2002; Duncan, 2010; Fedorenko et al., 2013b) across many domains, including language (Rodd et al., 2005; Novais-Santos et al., 2007; January et al., 2009; McMillan et al., 2012; Nieuwland et al., 2012; Wild et al., 2012; McMillan et al., 2013). An important goal for future work is thus to understand the division of labor between language and MD regions during syntactic processing (see also Fedorenko, 2014, for discussion). For example, which regions exhibit sensitivity to syntactic complexity earlier? Does activity in each network relate to distinct aspects of behavior? Is MD activity causally important for language comprehension (e.g., Amici et al., 2007)?

3.4.2. Is syntactic processing cognitively inseparable from other aspects of language comprehension?

Perhaps the most important consequence of the finding that syntactic processing is not localized to a subset of the language network is the suggestion of strong (and probably complete) overlap between regions that support syntactic processing and those that process word-level meanings (e.g., Fedorenko et al., 2012b; see Bates and Goodman, 1999, for an earlier extensive review and discussion; cf. Marin et al., 1976; Caramazza and Berndt, 1978, for earlier opposing views). Indeed, lexico-semantic processing appears to be similarly distributed across the language network. For example, contrasts between single words and various baselines (fixation, false fonts, pseudowords, etc.) elicit responses in all the language regions considered here (Humphries et al., 2007; Diaz and McCarthy, 2009; Fedorenko et al., 2010; Bedny et al., 2011).

Of course, it is not straightforward to compare roughly similar distributions of syntactic and lexico-semantic effects across separate studies, especially given the high inter-individual variability in the precise anatomical locations of language regions. It is possible that, within the same individual, each language region consists of several sub-regions, some more heavily recruited during syntactic processing and other more heavily recruited during lexico-semantic processing. Sub-regions of the latter kind might have been missed by our language localizer contrast (*sentences > nonwords*) if this contrast was somehow biased, such that syntactic differences across its two conditions were stronger than lexico-semantic differences.

However, even when we change our localizer contrast to a “purely” lexical comparison between word lists and nonword lists, the identified language regions show the critical syntactic complexity effect. More generally, other studies have directly contrasted lexical and syntactic manipulations and found overlapping activations. For example, Roder et al. (2002; see also Keller et al., 2001) examined syntactically complex vs. simpler sentences that were made up of real words vs. pseudowords. Inferior frontal and posterior temporal regions showed sensitivity to both manipulations: sentences composed of real words produced stronger responses than pseudoword sentences, and syntactically complex sentences produced stronger responses than syntactically simpler sentences. Thus, at least at the spatial scale of voxels measured with fMRI, syntactic and lexico-semantic processes appear to recruit the same set of regions distributed across the entire language network.

What are the theoretical implications of an overlap between syntactic processes and lexico-semantic processes at the level of their neural implementation? Specifically, does such overlap indicate that these processes are cognitively inseparable? This conjecture is in line with most current linguistic frameworks and the wealth of available psycholinguistic evidence. Specifically, when we know a language, we possess (i) a large but limited inventory of linguistic knowledge representations (e.g., words); and (ii) an ability to combine these stored knowledge representations to form a potentially infinite number of new meanings, i.e., a compositional capacity (Frege, 1914). Early proposals (e.g., Chomsky, 1965) linked lexico-semantic processing to the storage component of language (i.e., our lexicon), and syntactic processing – to its combinatorial component. However, over the last several decades, the nature of stored linguistic representations has evolved to allow for greater complexity, including information about how morphemes and words can combine with one another (Joshi et al., 1975; Bresnan, 1982; Schabes et al., 1988; Pollard and Sag, 1994; Chomsky, 1995; Goldberg, 1995; Bybee, 1998, 2010; Jackendoff, 2002, 2007; Culicover and Jackendoff, 2005). Consequently, many current proposals construe language knowledge as a continuum from the sounds of the language, to morphemes and words, to more complex units like words stored with the syntactic/semantic contexts in which they frequently occur (the degree of abstractness of these contexts varies depending on the details of the particular proposal). This view is supported by much experimental work showing that comprehenders appear to keep track of co-occurrences at different grain sizes, crossing the boundaries between words and combinatorial rules (Clifton et al., 1984; MacDonald et al., 1994; Trueswell et al., 1994; Garnsey et al., 1997; Traxler et al., 2002; Reali and Christiansen, 2007; Gennari and MacDonald, 2008), or between sounds and words (Farmer et al., 2006; Schmidtke et al., 2014). A similar picture obtains in the domain of language production (for a review, see Vigliocco and Hartsuiker, 2002).

Strong neuro-scientific support for the cognitive inseparability of syntactic and lexico-semantic processes cannot, however, rely on spatial overlap alone. It also requires (i) evidence for temporal overlap between the different processes recruiting a given language region; and (ii) causal evidence that the region in question is necessary for the different processes. Unfortunately, joint temporal, spatial and causal evidence cannot be obtained with fMRI. As discussed earlier, it requires methods such as electrocortical stimulation (Ojemann et al., 1989). Still, the spatial overlap between the responses to individual word meanings and to syntactic complexity throughout the language network allows us to at least entertain the hypothesis that the very same brain regions (i) store our language knowledge, and (ii) support the combination of those knowledge representations to form new meanings (see Hasson et al., 2015, for a recent discussion of this idea as applied to neural computation in general; cf. proposals like that of Baggio & Hagoort, 2001, according to which different brain regions of the language network support storage vs. combinatorial processing).

3.4.3. Dissociations within the language network?

As we argued in the Introduction, uncovering the division of linguistic labor among the regions of the fronto-temporal language network is key to understanding the cognitive architecture of the language faculty. However, the most fundamental aspects of the language network's architecture remain to be discovered. For example, how to divide the language network into constituent regions is still under debate: the division into eight regions based on the average topography of language activations adopted here (Fedorenko et al., 2010) is only a suggestion (see also Mahowald and Fedorenko, 2016). In fact, it is not even clear whether division of the language network into regions is warranted. On the one hand, the different regions of the language network show broadly similar functional profiles as measured with fMRI: they all respond more to meaningful and structured language stimuli like phrases and sentences than to "degraded" stimuli like lists of words, Jabberwocky sentences or lists of nonwords (Fedorenko et al., 2010; Baggio and Hagoort, 2011; Pallier et al., 2011). As shown here, they also all show sensitivity to finer-grain syntactic manipulations. In addition, language regions exhibit synchronized low-frequency oscillations during rest (Cordes et al., 2000; Hampson et al., 2002; Dronkers, 2011; Newman et al., 2013; Yue et al., 2013; Blank et al., 2014) and language comprehension (Blank et al., 2014). Finally, various functional properties of the language regions – e.g., how large or lateralized they are – are strongly correlated across regions (Mahowald and Fedorenko, 2016). All these results suggest that language regions form a functionally integrated network and should be considered as such when thinking about the architecture of language processing (e.g., Fedorenko and Thompson-Schill, 2014).

On the other hand, this is not to say that no functional dissociations exist within the language network. Indeed, a number of prior studies have reported differences among some of the language regions (Thompson-Schill et al., 1999; Bedny et al., 2008; Snijders et al., 2009; Mesulam et al., 2015). As discussed above, we should also keep in mind the low temporal resolution of fMRI: it is possible that dissociations would be more apparent when examining the language network through a finer temporal lens. Nevertheless, if one is to argue that some region or regions of the language network are functionally distinct from the rest of it, region by condition interactions are critical, and differences in overall responsiveness to language may need to be taken into account.

3.4.4. Conclusion

Our study provides evidence that sensitivity to syntactic complexity is widespread across the language network, contrary to many previous neuroimaging studies that reported only a few, localized foci of syntactic complexity effects. Investigations of syntactic processing therefore need to expand their scope to include the entire network of high-level language processing regions in order to fully understand how syntax is instantiated in the human brain. More generally, we recommend that neuroimaging studies of the language network follow two methodological considerations. First, analysis methods should allow for inter-individual variability in the exact anatomical location of functional regions. In this regard, functional localization of language regions individually in each participant is one promising method, showing increased sensitivity compared to traditional group analyses. Second, any hypothesized functional differences across regions of the language network should be tested by directly comparing effect sizes across regions (i.e., explicitly testing for a region-by-condition interaction), while taking into account more general differences in overall sensitivity to language. These considerations should guide us as we continue to accumulate evidence about the functional profiles of the regions of the language network; they will enable us to advance and evaluate specific hypotheses about the kinds of representations that such regions are likely to store and the computations that they are likely to perform.

3.5. Acknowledgements

We thank Nancy Kanwisher and Ted Gibson for many helpful discussions of this work. We thank Melissa Kline for recording the sentences. For comments on the earlier drafts of the manuscript we are grateful to Adele Goldberg, Marina Bedny, David Caplan, Nancy Kanwisher, and especially Ted Gibson, who read and commented on multiple versions. We would also like to acknowledge the Athinoula A. Martinos Imaging Center at McGovern Institute for Brain Research at MIT, and the support team (Steve Shannon, Atsushi Takahashi, and Sheeba Arnold). This work was supported by a K99/R00 award HD 057522 to E.F. from NICHD. K.M. was supported by a National Defense Science and Engineering Graduate (NDSEG) Fellowship.

References

Amici S, Ogar J, Brambati SM, Miller BL, Neuhaus J, Dronkers NL, Gorno-Tempini ML (2007) Performance in specific language tasks correlates with regional volume changes in progressive aphasia. Cognitive and Behavioral Neurology 20:203-211.

Amunts K, Schleicher A, Bürgel U, Mohlberg H, Uylings H, Zilles K (1999) Broca's region revisited: Cytoarchitecture and intersubject variability. Journal of Comparative Neurology 412:319-341.

Ances BM, Leontiev O, Perthen JE, Liang C, Lansing AE, Buxton RB (2008) Regional differences in the coupling of cerebral blood flow and oxygen metabolism changes in response to activation: Implications for BOLD-fMRI. NeuroImage 39:1510-1521.

Baggio G, Hagoort P (2011) The balance between memory and unification in semantics: A dynamic account of the N400. Language and Cognitive Processes 26:1338-1367.

Barde LH, Yeatman JD, Lee ES, Glover G, Feldman HM (2012) Differences in neural activation between preterm and full term born adolescents on a sentence comprehension task: Implications for educational accommodations. Developmental Cognitive Neuroscience 2:S114-S128.

Baron SG, Osherson D (2011) Evidence for conceptual combination in the left anterior temporal lobe. NeuroImage 55:1847-1852.

Basso A, Lecours AR, Moraschini S, Vanier M (1985) Anatomoclinical correlations of the aphasias as defined through computerized tomography: Exceptions. Brain and Language 26:201-229.

Bastiaanse R, Edwards S (2004) Word order and finiteness in Dutch and English Broca's and Wernicke's aphasia. Brain and Language 89:91-107.

Bates E, Goodman J (1999) On the emergence of language from the lexicon. In: The Emergence of Language (MacWhinney B, ed), pp 29-80. Mahwah, NJ: Lawrence Erlbaum.

Baudiffier V, Caplan D, Gaonac'h D, Chesnet D (2011) The effect of noun animacy on the processing of unambiguous sentences: Evidence from French relative clauses. The Quarterly Journal of Experimental Psychology 64:1896-1905.

Bedny M, Caramazza A, Grossman E, Pascual-Leone A, Saxe R (2008) Concepts are more than percepts: The case of action verbs. The Journal of Neuroscience 28:11347-11353.

Bedny M, Pascual-Leone A, Dodell-Feder D, Fedorenko E, Saxe R (2011) Language processing in the occipital cortex of congenitally blind adults. Proceedings of the National Academy of Sciences 108:4429-4434.

Bemis DK, Pylkkänen L (2011) Simple composition: A magnetoencephalography investigation into the comprehension of minimal linguistic phrases. The Journal of Neuroscience 31:2801-2814.

Ben-Shachar M, Hendler T, Kahn I, Ben-Bashat D, Grodzinsky Y (2003) The neural reality of syntactic transformations evidence from functional magnetic resonance imaging. Psychological Science 14:433-440.

Binder JR, Frost JA, Hammeke TA, Cox RW, Rao SM, Prieto T (1997) Human brain language areas identified by functional magnetic resonance imaging. The Journal of Neuroscience 17:353-362.

Blank I, Kanwisher N, Fedorenko E (2014) A functional dissociation between language and multiple-demand systems revealed in patterns of BOLD signal fluctuations. Journal of Neurophysiology 112:1105-1118.

Bornkessel I, Schlesewsky M (2006) The extended argument dependency model: A neurocognitive approach to sentence comprehension across languages. Psychological Review 113:787.

Bornkessel I, Zysset S, Friederici AD, von Cramon DY, Schlesewsky M (2005) Who did what to whom? The neural basis of argument hierarchies during language comprehension. NeuroImage 26:221-233.

Boston MF, Hale JT, Vasishth S, Kliegl R (2011) Parallel processing and sentence comprehension difficulty. Language and Cognitive Processes 26:301-349.

Brennan J, Pylkkänen L (2012) The time-course and spatial distribution of brain activity associated with sentence processing. NeuroImage 60:1139-1148.

Brennan J, Nir Y, Hasson U, Malach R, Heeger DJ, Pylkkänen L (2012) Syntactic structure building in the anterior temporal lobe during natural story listening. Brain and Language 120:163-173.

Bresnan J (1982) The Mental Representation of Grammatical Relations. Cambridge, MA: MIT Press.

Bybee J (1998) A functionalist approach to grammar and its evolution. Evolution of Communication 2:249-278.

Bybee J (2010) Language, Usage and Cognition. Cambridge, UK: Cambridge University Press.

Caplan D, Futter C (1986) Assignment of thematic roles to nouns in sentence comprehension by an agrammatic patient. Brain and Language 27:117-134.

Caplan D, Hildebrandt N, Makris N (1996) Location of lesions in stroke patients with deficits in syntactic processing in sentence comprehension. Brain 119:933-949.

Caplan D, Stanczak L, Waters G (2008) Syntactic and thematic constraint effects on blood oxygenation level dependent signal correlates of comprehension of relative clauses. Journal of Cognitive Neuroscience 20:643-656.

Caramazza A, Zurif EB (1976) Dissociation of algorithmic and heuristic processes in language comprehension: Evidence from aphasia. Brain and Language 3:572-582.

Caramazza A, Berndt RS (1978) Semantic and syntactic processes in aphasia: A review of the literature. Psychological Bulletin 85:898.

Chomsky N (1965) Aspects of the Theory of Syntax. Cambridge, MA: MIT press.

Chomsky N (1995) The Minimalist Program. Cambridge, MA: MIT Press.

Clifton C, Frazier L, Connine C (1984) Lexical expectations in sentence comprehension. Journal of Verbal Learning and Verbal Behavior 23:696-708.

Constable RT, Pugh KR, Berroya E, Mencl WE, Westerveld M, Ni W, Shankweiler D (2004) Sentence complexity and input modality effects in sentence comprehension: an fMRI study. NeuroImage 22:11-21.

Cooke A, Grossman M, DeVita C, Gonzalez-Atavales J, Moore P, Chen W, Gee J, Detre J (2006) Large-scale neural network for sentence processing. Brain and Language 96:14-36.

Cooke A, Zurif EB, DeVita C, Alsop D, Koenig P, Detre J, Gee J, Pinango M, Balogh J, Grossman M (2002) Neural basis for sentence comprehension: Grammatical and short - term memory components. Human Brain Mapping 15:80-94.

Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. Nature Reviews Neuroscience 3:201-215.

Cordes D, Haughton VM, Arfanakis K, Wendt GJ, Turski PA, Moritz CH, Quigley MA, Meyerand ME (2000) Mapping functionally related regions of brain with functional connectivity MR imaging. American Journal of Neuroradiology 21:1636-1644.

Culicover PW, Jackendoff R (2005) Simpler Syntax. Oxford, UK: Oxford University Press.

De Vries MH, Monaghan P, Knecht S, Zwitserlood P (2008) Syntactic structure and artificial grammar learning: The learnability of embedded hierarchical structures. Cognition 107:763-774.

Demberg V, Keller F (2008) Data from eye-tracking corpora as evidence for theories of syntactic processing complexity. Cognition 109:193-210.

Devlin JT, Russell RP, Davis MH, Price CJ, Wilson J, Moss HE, Matthews PM, Tyler LK (2000) Susceptibility-induced loss of signal: Comparing PET and fMRI on a semantic task. NeuroImage 11:589-600.

Diaz MT, McCarthy G (2009) A comparison of brain activity evoked by single content and function words: an fMRI investigation of implicit word processing. Brain Research 1282:38-49.

Dick F, Bates E, Wulfeck B, Utman JA, Dronkers N, Gernsbacher MA (2001) Language deficits, localization, and grammar: Evidence for a distributive model of language breakdown in aphasic patients and neurologically intact individuals. Psychological Review 108:759-788.

Dronkers NF (2011) The neural architecture of the language comprehension network: Converging evidence from lesion and connectivity analyses. Frontiers in Systems Neuroscience 5:1.

Dronkers NF, Ludy CA, Redfern BB (1998) Pragmatics in the absence of verbal language: Descriptions of a severe aphasic and a language-deprived adult. Journal of Neurolinguistics 11:179-190.

Dronkers NF, WIILKINS DP, VanValin RD, Redfern BB, Jaeger JJ (1994) A reconsideration of the brain-areas involved in the disruption of morphosyntactic comprehension. Brain and Language 47:461-463.

Dronkers NF, Wilkins DP, Van Valin RD, Redfern BB, Jaeger JJ (2004) Lesion analysis of the brain areas involved in language comprehension. Cognition 92:145-177.

Duffau H, Moritz-Gasser S, Mandonnet E (2014) A re-examination of neural basis of language processing: Proposal of a dynamic hodotopical model from data provided by brain stimulation mapping during picture naming. Brain and Language 131:1-10.

Duncan J (2010) The multiple-demand (MD) system of the primate brain: Mental programs for intelligent behaviour. Trends in Cognitive Sciences 14:172-179.

Duncan J, Owen AM (2000) Common regions of the human frontal lobe recruited by diverse cognitive demands. Trends in Neurosciences 23:475-483.

Ekstrom A (2010) How and when the fMRI BOLD signal relates to underlying neural activity: The danger in dissociation. Brain Research Reviews 62:233-244.

Embick D, Marantz A, Miyashita Y, O'Neil W, Sakai KL (2000) A syntactic specialization for Broca's area. Proceedings of the National Academy of Sciences 97:6150-6154.

Farmer TA, Christiansen MH, Monaghan P (2006) Phonological typicality influences on-line sentence comprehension. Proceedings of the National Academy of Sciences 103:12203-12208.

Fedorenko E (2014) The role of domain-general cognitive control in language comprehension. Frontiers in Psychology 5:335.

Fedorenko E, Kanwisher N (2011) Some regions within broca's area do respond more strongly to sentences than to linguistically degraded stimuli: A comment on Rogalsky and Hickok (2011). Journal of Cognitive Neuroscience 23:2632-2635.

Fedorenko E, Thompson-Schill SL (2014) Reworking the language network. Trends in Cognitive Sciences 18:120-126.

Fedorenko E, Gibson E, Rohde D (2006) The nature of working memory capacity in sentence comprehension: Evidence against domain-specific working memory resources. Journal of Memory and Language 54:541-553.

Fedorenko E, Behr MK, Kanwisher N (2011) Functional specificity for high-level linguistic processing in the human brain. Proceedings of the National Academy of Sciences 108:16428-16433.

Fedorenko E, Duncan J, Kanwisher N (2012a) Language-selective and domain-general regions lie side by side within Broca's area. Current Biology 22:2059-2062.

Fedorenko E, Nieto-Castañón A, Kanwisher N (2012b) Lexical and syntactic representations in the brain: an fMRI investigation with multi-voxel pattern analyses. Neuropsychologia 50:499-513.

Fedorenko E, Woodbury R, Gibson E (2013a) Direct evidence of memory retrieval as a source of difficulty in non - Local dependencies in language. Cognitive Science 37:378-394.

Fedorenko E, Duncan J, Kanwisher N (2013b) Broad domain generality in focal regions of frontal and parietal cortex. Proceedings of the National Academy of Sciences 110:16616-16621.

Fedorenko E, Hsieh P-J, Nieto-Castañón A, Whitfield-Gabrieli S, Kanwisher N (2010) New method for fMRI investigations of language: Defining ROIs functionally in individual subjects. Journal of Neurophysiology 104:1177-1194.

Fiebach CJ, Schlesewsky M, Friederici AD (2001) Syntactic working memory and the establishment of filler-gap dependencies: Insights from ERPs and fMRI. Journal of Psycholinguistic Research 30:321-338.

Fiebach CJ, Schlesewsky M, Lohmann G, Von Cramon DY, Friederici AD (2005) Revisiting the role of Broca's area in sentence processing: Syntactic integration versus syntactic working memory. Human Brain Mapping 24:79-91.

Fischl B, Rajendran N, Busa E, Augustinack J, Hinds O, Yeo BTT, Mohlberg H, Amunts K, Zilles K (2008) Cortical folding patterns and predicting cytoarchitecture. Cerebral Cortex 18:1973-1980.

Frazier L (1987) Syntactic processing: Evidence from Dutch. Natural Language and Linguistic Theory 5:519-559.

Frazier L, d'Arcais GBF (1989) Filler driven parsing: A study of gap filling in Dutch. Journal of Memory and Language 28:331-344.

Frege G (1914) Letter to Jourdain. In: Philosophical and Mathematical Correspondence (G. G, ed), pp 78-80. Chicago, IL: Chicago University Press.

Friederici AD (2009) Pathways to language: Fiber tracts in the human brain. Trends in Cognitive Sciences 13:175-181.

Friederici AD (2011) The brain basis of language processing: From structure to function. Physiological Reviews 91:1357-1392.

Friederici AD (2012) The cortical language circuit: From auditory perception to sentence comprehension. Trends in Cognitive Sciences 16:262-268.

Friederici AD, Kotz SA, Scott SK, Obleser J (2010) Disentangling syntax and intelligibility in auditory language comprehension. Human Brain Mapping 31:448-457.

Frost MA, Goebel R (2012) Measuring structural–functional correspondence: Spatial variability of specialised brain regions after macro-anatomical alignment. NeuroImage 59:1369-1381.

Garnsey SM, Pearlmutter NJ, Myers E, Lotocky MA (1997) The contributions of verb bias and plausibility to the comprehension of temporarily ambiguous sentences. Journal of Memory and Language 37:58-93.

Gennari SP, MacDonald MC (2008) Semantic indeterminacy in object relative clauses. Journal of Memory and Language 58:161-187.

Gennari SP, MacDonald MC (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.

Gibson E, Pearlmutter NJ (1998) Constraints on sentence comprehension. Trends in Cognitive Sciences 2:262-268.

Gibson E, Tily H, Fedorenko E (2013) The processing complexity of English relative clauses. In: Language Down the Garden Path: The Cognitive and Biological Basis for Linguistic Structure (Sanz M, Laka I, Tanenhaus MK, eds). Oxford, UK: Oxford University Press.

Goldberg AE (1995) Constructions: A construction grammar approach to argument structure. Chicago, IL: University of Chicago Press.

Gordon PC, Hendrick R, Johnson M (2001) Memory interference during language processing. Journal of Experimental Psychology: Learning, Memory, and Cognition 27:1411-1423.

Gordon PC, Hendrick R, Johnson M (2004) Effects of noun phrase type on sentence complexity. Journal of Memory and Language 51:97-114.

Grodner D, Gibson E (2005) Consequences of the serial nature of linguistic input for sentenial complexity. Cognitive Science 29:261-290.

Grodzinsky Y (2000) The neurology of syntax: Language use without Broca's area. Behavioral and Brain Sciences 23:1-21.

Grodzinsky Y, Friederici AD (2006) Neuroimaging of syntax and syntactic processing. Current Opinion in Neurobiology 16:240-246.

Grodzinsky Y, Santi A (2008) The battle for Broca's region. Trends in Cognitive Sciences 12:474-480.

Hale J (2001) A probabilistic earley parser as a psycholinguistic model. In: Proceedings of the second meeting of the North American Chapter of the Association for Computational Linguistics on Language technologies, pp 1-8: Association for Computational Linguistics.

Hampson M, Peterson BS, Skudlarski P, Gatenby JC, Gore JC (2002) Detection of functional connectivity using temporal correlations in MR images. Human Brain Mapping 15:247-262.

Harrison RV, Harel N, Panesar J, Mount RJ (2002) Blood capillary distribution correlates with hemodynamic-based functional imaging in cerebral cortex. Cerebral Cortex 12:225-233.

Hasson U, Chen J, Honey CJ (2015) Hierarchical process memory: Memory as an integral component of information processing. Trends in Cognitive Sciences 19:304-313.

Herrmann B, Obleser J, Kalberlah C, Haynes JD, Friederici AD (2012) Dissociable neural imprints of perception and grammar in auditory functional imaging. Human Brain Mapping 33:584-595.

Holmes VM, O'Regan JK (1981) Eye fixation patterns during the reading of relative-clause sentences. Journal of Verbal Learning and Verbal Behavior 20:417-430.

Humphries C, Willard K, Buchsbaum B, Hickok G (2001) Role of anterior temporal cortex in auditory sentence comprehension: an fMRI study. Neuroreport 12:1749-1752.

Humphries C, Love T, Swinney D, Hickok G (2005) Response of anterior temporal cortex to syntactic and prosodic manipulations during sentence processing. Human Brain Mapping 26:128-138.

Humphries C, Binder JR, Medler DA, Liebenthal E (2006) Syntactic and semantic modulation of neural activity during auditory sentence comprehension. Journal of Cognitive Neuroscience 18:665-679.

Humphries C, Binder JR, Medler DA, Liebenthal E (2007) Time course of semantic processes during sentence comprehension: an fMRI study. NeuroImage 36:924-932.

Ishizuka T, Nakatani K, Gibson E (2003) Relative clause extraction complexity in Japanese. In: Poster presented at the 16th annual CUNY conference on human sentence processing. Cambridge, MA.

Jackendoff R (1972) Semantic interpretation in generative grammar. Cambridge, MA: MIT Press.

Jackendoff R (2002) Foundation of Language: Brain, Meaning, Grammar, Evolution. Oxford, UK: Oxford University Press.

Jackendoff R (2007) A parallel architecture perspective on language processing. Brain Research 1146:2-22.

January D, Trueswell JC, Thompson-Schill SL (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 21:2434-2444.

Jefferies E (2013) The neural basis of semantic cognition: Converging evidence from neuropsychology, neuroimaging and TMS. Cortex 49:611-625.

Jezzard P, Clare S (1999) Sources of distortion in functional MRI data. Human Brain Mapping 8:80-85.

Joshi AK, Levy LS, Takahashi M (1975) Tree adjunct grammars. Journal of Computer and System Sciences 10:136-163.

Julian JB, Fedorenko E, Webster J, Kanwisher N (2012) An algorithmic method for functionally defining regions of interest in the ventral visual pathway. NeuroImage 60:2357-2364.

Just MA, Carpenter PA, Keller TA, Eddy WF, Thulborn KR (1996) Brain activation modulated by sentence comprehension. Science 274:114-116.

Kaan E, Swaab TY (2002) The brain circuitry of syntactic comprehension. Trends in Cognitive Sciences 6:350-356.

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 107:11163-11170.

Keller TA, Carpenter PA, Just MA (2001) The neural bases of sentence comprehension: a fMRI examination of syntactic and lexical processing. Cerebral Cortex 11:223-237.

Kriegeskorte N, Simmons WK, Bellgowan PSF, Baker CI (2009) Circular analysis in systems neuroscience: The dangers of double dipping. Nature Neuroscience 12:535-540.

Kwon N, Polinsky M, Kluender R (2006) Subject preference in Korean. In: Proceedings of the 25th West Coast Conference on Formal Linguistics, pp 1-14. Seattle, WA.

Kwon N, Gordon PC, Lee Y, Kluender R, Polinsky M (2010) Cognitive and linguistic factors affecting subject/object asymmetry: An eye-tracking study of prenominal relative clauses in Korean. Language 86:546-582.

Lambon-Ralph MA, Sage K, Jones RW, Mayberry EJ (2010) Coherent concepts are computed in the anterior temporal lobes. Proceedings of the National Academy of Sciences 107:2717-2722.

Levy RP (2008) Expectation-based syntactic comprehension. Cognition 106:1126-1177.

Levy RP, Fedorenko E, Gibson E (2013) The syntactic complexity of Russian relative clauses. Journal of Memory and Language 69:461-495.

Lewis RL, Vasishth S, Van Dyke JA (2006) Computational principles of working memory in sentence comprehension. Trends in Cognitive Sciences 10:447-454.

MacDonald MC, Pearlmutter NJ, Seidenberg MS (1994) The lexical nature of syntactic ambiguity resolution. Psychological Review 101:676-703.

Magnusdottir S, Fillmore P, Den Ouden DB, Hjaltason H, Rorden C, Kjartansson O, Bonilha L, Fridriksson J (2013) Damage to left anterior temporal cortex predicts impairment of complex syntactic processing: A lesion - symptom mapping study. Human Brain Mapping 34:2715-2723.

Mahowald K, Fedorenko E (2016) Reliable individual-level neural markers of high-level language processing: A necessary precursor for relating neural variability to behavioral and genetic variability. NeuroImage 139:74-93.

Mak WM, Vonk W, Schriefers H (2002) The influence of animacy on relative clause processing. Journal of Memory and Language 47:50-68.

Mak WM, Vonk W, Schriefers H (2006) Animacy in processing relative clauses: The hikers that rocks crush. Journal of Memory and Language 54:466-490.

Marin OS, Saffran EM, Schwartz MF (1976) Dissociations of language in aphasia: Implications for normal function. Annals of the New York Academy of Sciences 280:868-884.

Mazoyer BM, Tzourio N, Frak V, Syrota A, Murayama N, Levrier O, Salamon G, Dehaene S, Cohen L, Mehler J (1993) The cortical representation of speech. Journal of Cognitive Neuroscience 5:467-479.

McElree B, Foraker S, Dyer L (2003) Memory structures that subserve sentence comprehension. Journal of Memory and Language 48:67-91.

McMillan CT, Clark R, Gunawardena D, Ryant N, Grossman M (2012) fMRI evidence for strategic decision-making during resolution of pronoun reference. Neuropsychologia 50:674-687.

McMillan CT, Coleman D, Clark R, Liang TW, Gross RG, Grossman M (2013) Converging evidence for the processing costs associated with ambiguous quantifier comprehension. Frontiers in Psychology 4:153.

Mecklinger A, Schriefers H, Steinhauer K, Friederici AD (1995) Processing relative clauses varying on syntactic and semantic dimensions - an analysis with event-related potentials. Memory and Cognition 23:477-494.

Meltzer JA, McArdle JJ, Schafer RJ, Braun AR (2010) Neural aspects of sentence comprehension: Syntactic complexity, reversibility, and reanalysis. Cerebral Cortex 20:1853-1864.

Menenti L, Segaert K, Hagoort P (2012) The neuronal infrastructure of speaking. Brain and Language 122:71-80.

Menon RS, Kim SG (1999) Spatial and temporal limits in cognitive neuroimaging with fMRI. Trends in Cognitive Sciences 3:207-216.

Mesulam MM (1990) Large-scale neurocognitive networks and distributed-processing for attention, language, and memory. Ann Neurol 28:597-613.

Mesulam MM, Thompson CK, Weintraub S, Rogalski EJ (2015) The Wernicke conundrum and the anatomy of language comprehension in primary progressive aphasia. Brain 138:2423-2437.

Miyamoto ET, Nakamura M (2003) Subject/object asymmetries in the processing of relative clauses in Japanese. In: Proceedings of the 22nd West Coast Conference on Formal Linguistics, pp 342-355. San Diego, CA.

Monti MM, Parsons LM, Osherson DN (2012) Thought beyond language: neural dissociation of algebra and natural language. Psychological Science 23:914-922.

Newman AJ, Kenny S, Saint-Aubin J, Klein RM (2013) Can skilled readers perform a second task in parallel? A functional connectivity MRI study. Brain and Language 124:84-95.

Nieto-Castañón A, Fedorenko E (2012) Subject-specific functional localizers increase sensitivity and functional resolution of multi-subject analyses. NeuroImage 63:1646-1669.

Nieuwenhuis S, Forstmann BU, Wagenmakers E-J (2011) Erroneous analyses of interactions in neuroscience: A problem of significance. Nature Neuroscience 14:1105-1107.

Nieuwland MS, Martin AE, Carreiras M (2012) Brain regions that process case: Evidence from Basque. Human Brain Mapping 33:2509-2520.

Novais-Santos S, Gee J, Shah M, Troiani V, Work M, Grossman M (2007) Resolving sentence ambiguity with planning and working memory resources: Evidence from fMRI. NeuroImage 37:361-378.

O'Grady W, Lee M, Choo M (2003) A subject-object asymmetry in the acquisition of relative clauses in Korean as a second language. Studies in Second Language Acquisition 25:433-448.

O'Grady W (2011) Relative clauses: Processing and acquisition. In: The Acquisition of Relative Clauses: Processing, Typology and Function (Kidd E, ed), pp 13-38. Philadelphia, PA: John Benjamin Publishing.

Ojemann G, Ojemann J, Lettich E, Berger M (1989) Cortical language localization in left, dominant hemisphere: An electrical stimulation mapping investigation in 117 patients. Journal of Neurosurgery 71:316-326.

Pallier C, Devauchelle A-D, Dehaene S (2011) Cortical representation of the constituent structure of sentences. Proceedings of the National Academy of Sciences 108:2522-2527.

Papoutsi M, Stamatakis EA, Griffiths J, Marslen-Wilson WD, Tyler LK (2011) Is left fronto-temporal connectivity essential for syntax? Effective connectivity, tractography and performance in left-hemisphere damaged patients. NeuroImage 58:656-664.

Pascual B, Masdeu JC, Hollenbeck M, Makris N, Insausti R, Ding SL, Dickerson BC (2015) Large-scale brain networks of the human left temporal pole: A functional connectivity MRI study. Cerebral Cortex 25:680-702.

Patterson K, Nestor PJ, Rogers TT (2007) Where do you know what you know? The representation of semantic knowledge in the human brain. Nature Reviews Neuroscience 8:976-987.

Peelle JE, Troiani V, Wingfield A, Grossman M (2010) Neural processing during older adults' comprehension of spoken sentences: Age differences in resource allocation and connectivity. Cerebral Cortex 20:773-782.

Pollard C, Sag IA (1994) Head-driven phrase structure grammar. Chicago, IL: University of Chicago Press.

Reali F, Christiansen MH (2007) Processing of relative clauses is made easier by frequency of occurrence. Journal of Memory and Language 57:1-23.

Rodd JM, Davis MH, Johnsrude IS (2005) The neural mechanisms of speech comprehension: fMRI studies of semantic ambiguity. Cerebral Cortex 15:1261-1269.

Roder B, Stock O, Neville H, Bien S, Rosler F (2002) Brain activation modulated by the comprehension of normal and pseudo-word sentences of different processing demands: A functional magnetic resonance Imaging study. NeuroImage 15:1003-1014.

Rogalsky C, Hickok G (2009) Selective attention to semantic and syntactic features modulates sentence processing networks in anterior temporal cortex. Cerebral Cortex 19:786-796.

Rogalsky C, Hickok G (2011) The role of Broca's area in sentence comprehension. Journal of Cognitive Neuroscience 23:1664-1680.

Rolheiser T, Stamatakis EA, Tyler LK (2011) Dynamic processing in the human language system: Synergy between the arcuate fascicle and extreme capsule. J Neurosci 31:16949-16957.

Samuels JA, Benson DF (1979) Some aspects of language comprehension in anterior aphasia. Brain and Language 8:275-286.

Santi A, Grodzinsky Y (2010) fMRI adaptation dissociates syntactic complexity dimensions. NeuroImage 51:1285-1293.

Saxe R, Brett M, Kanwisher N (2006) Divide and conquer: A defense of functional localizers. NeuroImage 30:1088-1096.

Schabes Y, Abeille A, Joshi AK (1988) Parsing strategies with "lexicalized" grammars: Application to tree adjoining grammars. In: Proceedings of the 12th Conference on Computational Linguistics, pp 578-583. Stroudsburg, PA.

Schlesewsky M, Fanselow G, Kliegl R, Krems J (2000) The subject preference in the processing of locally ambiguous wh-questions in German. In: German Sentence Processing (Hemforth B, Konieczny L, eds), pp 65-93. Dordrecht, The Netherlands: Springer.

Schmidtke DS, Conrad M, Jacobs AM (2014) Phonological iconicity. Frontiers in Psychology 5:80.

Schriefers H, Friederici AD, Kuhn K (1995) The processing of locally ambiguous relative clauses in German. Journal of Memory and Language 34:499-520.

Schwartz MF, Saffran EM, Marin OSM (1980) The word order problem in agrammatism 1: Comprehension. Brain and Language 10:249-262.

Segaert K, Menenti L, Weber K, Petersson KM, Hagoort P (2012) Shared syntax in language production and language comprehension: an fMRI Study. Cerebral Cortex 22:1662-1670.

Selnes OA, Knopman DS, Niccum N, Rubens AB, Larson D (1983) Computed tomographic scan correlates of auditory comprehension deficits in aphasia: A prospective recovery study. Ann Neurol 13:558-566.

Snijders TM, Vosse T, Kempen G, Van Berkum JJA, Petersson KM, Hagoort P (2009) Retrieval and unification of syntactic structure in sentence comprehension: an fMRI study using word-category ambiguity. Cerebral Cortex 19:1493-1503.

Stromswold K, Caplan D, Alpert N, Rauch S (1996) Localization of syntactic comprehension by positron emission tomography. Brain and Language 52:452-473.

Tahmasebi AM, Davis MH, Wild CJ, Rodd JM, Hakyemez H, Abolmaesumi P, Johnsrude IS (2012) Is the link between anatomical structure and function equally strong at all cognitive levels of processing? Cerebral Cortex 22:1593-1603.

Tanenhaus MK, Trueswell JC (1995) Sentence comprehension. In: Speech, Language and Communication (Miller JL, Eimas PD, eds), pp 217-262. San Diego, CA: Academic Press.

Thompson-Schill SL, D'Esposito M, Kan IP (1999) Effects of repetition and competition on activity in left prefrontal cortex during word generation. Neuron 23:513-522.

Thothathiri M, Kimberg DY, Schwartz MF (2012) The neural basis of reversible sentence comprehension: Evidence from voxel-based lesion symptom mapping in aphasia. Journal of Cognitive Neuroscience 24:212-222.

Tramo MJ, Baynes K, Volpe BT (1988) Impaired syntactic comprehension and production in Broca's aphasia - CT lesion localization and recovery patterns. Neurology 38:95-98.

Traxler MJ, Morris RK, Seely RE (2002) Processing subject and object relative clauses: Evidence from eye movements. Journal of Memory and Language 47:69-90.

Trueswell JC, Tanenhaus MK, Garnsey SM (1994) Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. Journal of Memory and Language 33:285-318.

Tyler LK, Marslen-Wilson WD, Randall B, Wright P, Devereux BJ, Zhuang J, Papoutsi M, Stamatakis EA (2011) Left inferior frontal cortex and syntax: Function, structure and behaviour in patients with left hemisphere damage. Brain 134:415-431.

Ueno M, Garnsey SM (2008) An ERP study of the processing of subject and object relative clauses in Japanese. Language and Cognitive Processes 23:646-688.

Ullman MT (2004) Contributions of memory circuits to language: The declarative/procedural model. Cognition 92:231-270.

Vandenberghe R, Nobre AC, Price CJ (2002) The response of left temporal cortex to sentences. Journal of Cognitive Neuroscience 14:550-560.

Varley RA, Klessinger NJC, Romanowski CAJ, Siegal M (2005) Agrammatic but numerate. Proceedings of the National Academy of Sciences 102:3519-3524.

Vigliocco G, Hartsuiker RJ (2002) The interplay of meaning, sound, and syntax in sentence production. Psychological Bulletin 128:442-472.

Wanner E, Maratsos M (1978) An ATN approach to comprehension. In: Linguistic Theory and Psychological Reality (M. H, Bresnan J, Miller G, eds), pp 119-161. Cambridge, MA: MIT Press.

Wartenburger I, Heekeren HR, Burchert F, De Bleser R, Villringer A (2003) Grammaticality judgments on sentences with and without movement of phrasal constituents: an event-related fMRI study. Journal of Neurolinguistics 16:301-314.

Wells JB, Christiansen MH, Race DS, Acheson DJ, MacDonald MC (2009) Experience and sentence processing: Statistical learning and relative clause comprehension. Cognitive Psychol 58:250-271.

Wild CJ, Yusuf A, Wilson DE, Peelle JE, Davis MH, Johnsrude IS (2012) Effortful listening: The processing of degraded speech depends critically on attention. J Neurosci 32:14010-14021.

Willems RM, Benn Y, Hagoort P, Toni I, Varley R (2011) Communicating without a functioning language system: Implications for the role of language in mentalizing. Neuropsychologia 49:3130-3135.

Wilson SM (2014) The impact of vascular factors on language localization in the superior temporal sulcus. Human Brain Mapping 35:4049-4063.

Wilson SM, Saygin AP (2004) Grammaticality judgment in aphasia: Deficits are not specific to syntactic structures, aphasic syndromes, or lesion sites. Journal of Cognitive Neuroscience 16:238-252.

Wilson SM, Galantucci S, Tartaglia MC, Rising K, Patterson DK, Henry ML, Ogar JM, DeLeon J, Miller BL, Gorno-Tempini ML (2011) Syntactic processing depends on dorsal language tracts. Neuron 72:397-403.

Wong C, Gallate J (2012) The function of the anterior temporal lobe: A review of the empirical evidence. Brain Research 1449:94-116.

Yue Q, Zhang L, Xu G, Shu H, Li P (2013) Task-modulated activation and functional connectivity of the temporal and frontal areas during speech comprehension. Neuroscience 237:87-95.

Zhang LM, Pylkkanen L (2015) The interplay of composition and concept specificity in the left anterior temporal lobe: An MEG study. NeuroImage 111:228-240.

Zurif E, Swinney D, Prather P, Solomon J, Bushell C (1993) An on-line analysis of syntactic processing in Broca's and Wernicke's aphasia. Brain and Language 45:448-464.

Appendix 3A: Sensitivity to syntactic complexity in the extended language network

Table 3A Effects[a] for the localizer contrast (*sentences* > *nonwords*)[b] and for the critical contrast (*object-extracted* > *subject-extracted*) in the extended language network

fROI	Localizer effect	Syntactic complexity effect
Right hemisphere homologues of the language network		
IFGorb	t=3.71; p<0.005	t<1; n.s.
IFG	t=3.89; p<0.005	t=2.19; p<0.05
MFG	t=2.05; p<0.05	t=1.23; n.s.
AntTemp	t=5.37; $p<10^{-4}$	t<1; n.s.
MidAntTemp	t=3.75; p<0.005	t<1; n.s.
MidPostTemp	t=5.67; $p<10^{-4}$	t=2.39; p<0.05
PostTemp	t=4.44; $p<10^{-3}$	t<1; n.s.
AngG	t=2.59; p<0.05	t<1; n.s.
Medial frontal cortex region		
Left SFG	t=5.51; $p<10^{-4}$	t<1; n.s.
Cerebellar regions		
Right Cereb	t=4.63; $p<10^{-3}$	t=1.88; p<0.05
Left Cereb	t=6.28; p<10^{-4}	t<1; n.s.

[a] We report uncorrected p values (df=12).

[b] Estimated in data not used for defining the fROIs, using across-runs cross-validation.

Chapter 4

A new functional signature of high-level language regions: common timescales for integrating information

With Evelina Fedorenko

Abstract

The cortical language network consists of left temporal and frontal regions that are selectively engaged in linguistic processing. How these regions functionally differ in their respective contributions to language comprehension has long been debated. Recently, a cortical hierarchy has been suggested in which linguistic input is integrated over increasingly longer time-windows, extending from mid-posterior/anterior temporal areas, through surrounding areas to inferior frontal areas (Lerner et al., 2011). Whereas the topography of this hierarchy appears to overlap with the language network, such correspondence is difficult to establish because the frontal and temporal lobes are functionally heterogeneous and the mapping of function onto their anatomy varies across individuals. Therefore, here we characterized the integration time-windows of seven language regions that were localized individually in each participant. Regions with short time-windows were expected to reliably track any locally well-structured input, unable to detect coarser, global incoherence, whereas regions with longer time-windows were expected to be sensitive to such coarse incoherence. Thus, participants listened to a story and three scrambled versions, shuffled at the scale of paragraphs, sentences, or words. The reliability of regional activity time-courses elicited by each stimulus was measured across participants (Hasson et al., 2004). Despite sufficient statistical power, no region-by-condition interactions were observed. Overall, language regions tracked both the intact story and scrambled paragraphs with high reliability, but reliability somewhat decreased for scrambled sentences and even more so for scrambled words. Therefore, the language network constitutes a unique stage of temporal integration within a broader cortical hierarchy.

4.1. Introduction

Language comprehension engages a cortical network of frontal and temporal brain regions, primarily in the left hemisphere (Bates et al., 2003; Fedorenko et al., 2010; Menenti et al., 2011). There is ample evidence that this network is language-selective and is not recruited by other cognitive processes (Fedorenko and Varley, 2016), but the division of linguistic labor among its constituent regions remains unclear. Whereas some neuroimaging studies suggest that different linguistic processes are localized to focal and distinct subsets of the language network (Dapretto and Bookheimer, 1999; Hickok and Poeppel, 2004; Ullman, 2004; Humphries et al., 2006; Hagoort, 2013), others report that different processes recruit widely distributed regions that spatially overlap (Keller et al., 2001; Fedorenko et al., 2012b; Bautista and Wilson, 2016; Blank et al., 2016). Similar inconsistencies plague the neuropsychological literature (Bates et al., 2003; Mesulam et al., 2015; Mirman et al., 2015).

Many studies search for a division of labor across processes that operate on linguistic units of different sizes: e.g., sub-lexical units (sounds), morphemes/words, and sentences (Mazoyer et al., 1993; Friederici, 2002; Price, 2012). Perhaps then linguistic processes can be differentiated by the size of the temporal window over which they integrate information: some processes have short integration time-windows and process any locally well-structured input, unable to detect coarser, global disorder (e.g., phonemes can be identified even in a scrambled sentence). Other processes integrate input over a longer time-scale sensitive to a broader context. However, current psycholinguistic theories and evidence suggest that language processing operates over a continuum of representations that lacks the traditionally postulated boundaries between sounds, words and grammar (Joshi et al., 1975; Bresnan, 1982; Schabes et al., 1988; Pollard and Sag, 1994; Chomsky, 1995; Goldberg, 1995; Bybee, 1998, 2010; Jackendoff, 2002, 2007; Culicover and Jackendoff, 2005). Therefore, linguistic labor might not be divided based on integration time-windows.

To test whether language-responsive cortex is functionally organized according to the length of integration time-windows, Lerner et al. (2011) recently measured how reliably different brain regions process linguistic input scrambled at different grain levels. Specifically, regional fluctuations in the fMRI BOLD signal were measured while individuals listened to a story or to increasingly scrambled variants consisting of reordered paragraphs, sentences, or words, and signal reliability was evaluated by comparing time-courses across individuals. This novel approach revealed a hierarchy of integration time-windows across the cortex, extending from mid-posterior and mid-anterior temporal regions both rostrally and caudally along the temporal lobe and on to inferior frontal regions. Early in the hierarchy, regions reliably tracked even finely-scrambled input, exhibiting a short integration time-window. Downstream regions became sensitive to increasingly coarser levels of scrambling, having increasingly longer integration time-windows.

Different stages of this hierarchy appear to anatomically overlap with different high-level language regions. Nevertheless, establishing a functional correspondence between this hierarchy and the language network is challenging. First, the topography of the temporal integration hierarchy was mapped by measuring the inter-individual reliability of signal time-courses on a voxel-by-voxel basis. However, function-to-voxel mapping in the temporal and frontal lobes exhibits high inter-individual variability so a given voxel might be part of the language network in one individual but be part of a different network in another (Gloor, 1997; Amunts et al., 1999; Wise et al., 2001; Poldrack, 2006; Fischl et al., 2008; Fedorenko et al., 2012a; Tahmasebi et al., 2012). Second, functional profiles across different stages of the hierarchy were not directly compared to one another, so it is unclear which regions meet the statistical standard for inferring a functional dissociation (Nieuwenhuis et al., 2011).

Therefore, here we explicitly characterized integration time-windows across the language network. To this end, we first functionally localized the language network in each individual participant (Fedorenko et al., 2010), thus accounting for inter-individual variability in its precise topography. Then, we examined the integration time-window of each region using the stimuli designed by Lerner et al. (2011). Finally, we directly compared the functional profiles across language regions to test whether their respective integration time-windows differed.

4.2. Materials and Methods

4.2.1. Participants

Seventeen participants (10 females) between the ages of 18 and 47, recruited from the MIT student body and the surrounding community, were paid for participation. All participants were native English speakers, had normal hearing, and gave informed consent in accordance with the requirements of MIT's Committee on the Use of Humans as Experimental Subjects (COUHES).

4.2.2. Data acquisition

Structural and functional data were collected on a whole-body 3 Tesla Siemens Trio scanner with a 32-channel head coil at the Athinoula A. Martinos Imaging Center at the McGovern Institute for

Brain Research at MIT. T1-weighted structural images were collected in 176 axial slices with 1mm isotropic voxels (repetition time (TR) = 2,530ms; echo time (TE) = 3.48ms). Functional, blood oxygenation level-dependent (BOLD) data were acquired using an EPI sequence with a 90° flip angle and using GRAPPA with an acceleration factor of 2; the following parameters were used: thirty-one 4.4mm thick near-axial slices acquired in an interleaved order (with 10% distance factor), with an in-plane resolution of 2.1×2.1mm, FoV in the phase encoding (A >> P) direction 200mm and matrix size 96x96mm, TR = 2000ms and TE = 30ms. The first 10s of each run were excluded to allow for steady state magnetization.

4.2.3. Design, materials and procedure

Each participant performed the language localizer task (Fedorenko et al., 2010) and, for the critical experiment, listened to five variants of a narrated story (cf. Lerner et al., 2011, where different groups listened to different variants). The localizer and critical experiment were run either in the same scanning session (11 participants) or in two separate sessions (6 participants, who have previously performed the localizer task while participating in other studies; see Mahowald & Fedorenko, in press, for evidence of high stability of language localizer activations over time). In each session, participants performed a few other, unrelated tasks, with scanning sessions lasting approximately 2h.

4.2.3.1. Language localizer task

Regions in the language network were localized using a passive reading task that contrasted sentences and lists of unconnected, pronounceable nonwords (Fedorenko et al., 2010). Each stimulus consisted of 12 words/nonwords, presented at the center of the screen one word/nonword at a time at a rate of 450ms per word/nonword. Each trial began with 100ms of fixation and ended with an icon instructing participants to press a button, presented for 400ms and followed by 100ms of fixation, for a total trial duration of 6s. The button-press task was included to help participants remain alert and focused throughout the run. Trials were presented in a standard blocked design with a counterbalanced order across two runs. Each block, consisting of 3 trials, lasted 18s. Fixation blocks were evenly distributed throughout the run and lasted 14s. Each run consisted of 8 blocks per condition and 5 fixation blocks, lasting a total of 358s. This localizer has been extensively validated, with language regions showing robust and reliable *sentences > nonwords* effects that generalize across materials, tasks (passive reading or memory) (Fedorenko et al., 2010) and modality of presentation (visual or auditory) (Braze et al., 2011; Vagharchakian et al., 2012; Fedorenko, 2014).

4.2.3.2. Critical experiment

Participants listened to the same materials that were originally used to characterize the cortical hierarchy of temporal integration windows (Lerner et al., 2011). These materials were based on an audio recording of a narrated story ("Pie-Main", told by Jim O'Grady at an event of "The Moth" group, NYC). They included: (i) the intact audio; (ii) three "scrambled" versions of the story that differed in the temporal scale of scrambling and were created by randomly reordering paragraphs, sentences or words, respectively; and (iii) an audio-reversed version. The last stimulus served as a low-level control condition, because reverse speech contains the same acoustic characteristics as speech and is similarly processed (Kimura and Folb, 1968; Koeda et al., 2006; but see Norman-Haignere et al., 2015), but does not carry linguistic information beyond the phonetic level (Binder et al., 2000). Therefore, regions with integrations time-windows that are sensitive to structure at the syllable level and above are not expected to reliably process this stimulus.

To render these materials suitable for our existing scanning protocol with a TR of 2s, the 15s of silence and/or music preceding and following each stimulus were each extended by 1s. The resulting 16s periods were not included in the analyses reported below. In addition, the two longest paragraphs in the reordered-paragraphs stimulus were each split into two sections, and one section was randomly repositioned in the stream of shuffled paragraphs. No other edits were made to the original materials.

Participants listened to the materials over scanner-safe headphones (Sensimetrics, Malden, MA), in one of two orders: for 10 participants, the intact story was played first and was followed by increasingly finer levels of temporal scrambling (from paragraphs to sentences to words). For the remaining 7 participants, the reordered-words stimulus was played first and was followed by decreasing levels of temporal scrambling (from sentences to paragraphs to the intact story). The reverse story was positioned either in the middle of the scanning session or at the end.

At the end of the scanning session, participants answered 8 multiple-choice questions concerning characters, places and events from particular points in the narrative, with foils describing information presented elsewhere in the story. All participants demonstrated good comprehension of the story (16 of them answered all questions correctly, and the remaining participant had only one error). Therefore, no participants were excluded from data analysis.

4.2.4. Data preprocessing

Preprocessing was performed using SPM5 and custom MATLAB scripts. Anatomical data were normalized into a common space (Montreal Neurological Institute (MNI) template), resampled into 2mm isotropic voxels, and segmented into probabilistic maps of the gray matter, white matter (WM) and cerebrospinal fluid (CSF). Functional data were motion corrected, resampled into 2mm isotropic voxels, smoothed with a 4mm FWHM Gaussian filter and high-pass filtered at 200s.

Additional temporal preprocessing of data from runs of the critical experiments was performed using the CONN toolbox (Whitfield-Gabrieli and Nieto-Castañón, 2012) with default parameters, unless specified otherwise. Five temporal principal components of the BOLD signal time-courses extracted from the WM were regressed out of each voxel's time-course; signal originating in the CSF was similarly regressed out. Six principal components of the six motion parameters estimated during offline motion correction were also regressed out, as well as their first time derivative. Next, the residual signal was bandpass filtered (0.008–0.09 Hz) to preserve only low-frequency signal fluctuations, because higher frequencies are contaminated by fluctuations originating from non-neural sources (Cordes et al., 2001). Nonetheless, the pattern of results we report below was fully obtained when we instead analyzed unfiltered time-courses that contained higher frequencies; however, these time-courses exhibited overall lower reliability across participants (reflecting the effect of idiosyncratic, non-neural fluctuations). We report the analyses of filtered data because they exhibit a wider dynamic range of response reliability across stimulus conditions and thus provide more power for comparisons across conditions and regions.

4.2.5. Data analysis

4.2.5.1. Functionally defining language regions in individual participants

Data from the language localizer task were analyzed using a general linear model that estimated the voxel-wise effect size of each condition in each experimental run. These effects were each modeled with a boxcar function (representing entire blocks) convolved with the canonical hemodynamic response function. The model also included first-order temporal derivatives of these effects, as well as nuisance regressors representing entire experimental runs and offline-estimated motion parameters. The obtained beta weights were then used to compute the voxel-wise *sentences > nonwords* contrast. The resulting contrast maps were restricted to include only gray matter voxels, excluding voxels that were more likely to belong to either the WM or the CSF based on segmentation of the participant's structural data.

Functional regions of interest (fROIs) in the language network were then defined using group-constrained, participant-specific localization (Fedorenko et al., 2010). For each participant, the map of the *sentences* > *nonwords* contrast was intersected with binary masks that constrained the participant-specific language network to fall within areas where activations for this contrast are relatively likely across the population. These masks are based on a group-level representation of the contrast obtained from a previous sample. We used 7 such masks, including left hemispheric regions in the posterior, mid-posterior, mid-anterior and anterior temporal lobe, as well as in the middle frontal gyrus, the inferior frontal gyrus and its orbital part (Figure 4.1). Critically, these masks cover large portions of the cortical hierarchy of temporal integration, including regions that have been suggested to differ from one another in their integration time-windows (Lerner et al., 2011). In each of these masks, a participant-specific fROI was defined as the top 10% of voxels with the highest contrast values. This top *n%* approach ensures that fROIs can be defined in every participant and that their sizes are the same across participants, allowing for generalizable results (Nieto-Castañón and Fedorenko, 2012).

Figure 4.1. Language fROIs for each of 3 sample participants. fROIs (red) are shown superimposed on an inflated brain. They were localized with a reading task, and were constrained to fall within seven broad areas where activations for this task are common across the population (pink).

4.2.5.2. Inter-subject correlation analysis

For each condition in the critical experiment, BOLD signal time-courses were extracted from each fROI in each participant. We analyzed time-courses beginning 6 seconds following the onset of the stimulus, in order to exclude an initial rise in the hemodynamic response relative to fixation; this rise is a trivially reliable component of the BOLD signal and might therefore blur differences in the reliability of signal fluctuations across different conditions and fROIs. Regional time-courses were computed by first *z*-scoring the time-course of each voxel and then averaging time-courses across voxels. Next, for each fROI, we computed Pearson's moment correlation coefficient between the time-course of one participant and the corresponding average time-course across the remaining participants. This procedure was iterated over all partitions of the participant pool, resulting in 17 ISCs per fROI. Correlations were Fisher-transformed to improve the normality of their distribution (Silver and Dunlap, 1987). These regional inter-subject correlations (ISCs) quantify how similar BOLD signal fluctuations are across participants, with high values indicative of regional activity that reliably tracks the incoming input (correlations across participants mirror correlations within a single participant across stimulus presentations; Golland et al., 2007; Hasson et al., 2009). Reliable tracking is expected only when the input is well structured at the temporal scale over which a fROI integrates information; weaker tracking is expected for inputs that are scrambled at that scale and thus cannot be reliably integrated. Therefore, comparing ISCs of a fROI across stimuli that vary in their level of scrambling serves to characterize its integration time-window.

Replicating the approach of Lerner et al. (2011), ISCs were tested for significance against an empirical null distribution (Theiler et al., 1992) and the resulting p-values were corrected for multiple comparisons using a false discovery rate (FDR) correction (Benjamini and Yekutieli, 2001). In addition, we directly tested whether ISCs differed across conditions. First, within each fROI, ISCs were compared across pairs of conditions using dependent-samples t-tests (FDR-corrected across all fROIs). Second, for the main analysis, we explicitly tested for a region-by-condition interaction using a two-way, repeated-measures analysis of variance (ANOVA) with fROI (7 levels) and condition (5 levels) as within-participant factors. To interpret our findings, we performed additional two-way ANOVAs that included all 7 fROIs but were each restricted to a specific pair of conditions, consisting of one scrambled stimulus and the intact story (FDR-corrected across condition pairs). Within each such test, we further performed 2×2 ANOVAs that were each restricted to a pair of fROIs (FDR-corrected across fROI pairs).

We obtained similar results when all tests were instead run as linear, mixed-effects regression analyses with fixed effects for fROI, condition and their interaction, as well as a random, by-participant intercept and random, by-participant slopes for both fROI and condition (Barr et al., 2013). Results also remained unaltered when we used non-parametric tests, namely, analytic permutation tests for 2×2 ANOVAs (Gill, 2007; Mewhort et al., 2010) and empirical permutations of reduced residuals for more complex ANOVAs (Anderson and Braak, 2003). Therefore, our findings are independent of assumptions regarding the distribution of ISCs.

4.3. Results

4.3.1. A lenient statistical approach reveals a temporal integration hierarchy

To characterize the integration time-windows of different language regions, we first tested which stimuli elicited significant ISCs in each of the 7 fROIs. This approach does not directly compare ISCs across conditions and, critically, does not test for a region-by-condition interaction. Nonetheless, we report this analysis in order to contrast it with the results obtained using more rigorous statistical tests, described below.

This analysis revealed that three fROIs – in the mid-posterior, mid-anterior and anterior temporal cortex – had significant ISCs for all five conditions, including the reverse story. If we were to determine their integration time-window based on these tests alone, we would conclude that they have an extremely short time-window, because their activity tracks even inputs that are ill-structured at the fine scale of participants' native (English) morphology. In the original cortical hierarchy of temporal integration, this time-window was associated with perceptual and sensory processes, yet our fROIs by definition extend well beyond low-level auditory areas (they respond more to sentences than to nonwords in both auditory and visual presentations, unlike sensory regions). Next, we found that frontal fROIs – in the middle frontal gyrus (MFG), inferior frontal gyrus (IFG) and its orbital part – significantly tracked all stimuli except for the reverse story. By the standards of the current analysis, their activity tracks inputs as long as they are well structured at the scale of morphemes or single words and, hence, that is their integration time-window. Finally, the posterior temporal fROI tracked the intact story, scrambled paragraphs and scrambled sentences but not the scrambled words or the reverse story; it would be characterized by a relatively long integration time-window, at the scale of sentences (Figure 4.2).

We note that, even with such a lenient statistical approach, we find that neural activity in many fROIs can tolerate higher levels of input scrambling than has previously been reported (Lerner et al., 2011; Friederici, 2012), suggestive of shorter integration windows. Specifically, significant ISCs for the reverse story might appear surprising. They are not likely to be caused by the processing of prosody, which is primarily processed by the right hemisphere (Fedorenko et al., 2015). Instead, they might reflect reliable processing of "foreign speech"-like information, or brief, low-level transitions between speech and silence. Regardless of these explanations, the power to detect such significant ISCs for relatively incoherent stimuli likely results from (i) a repeated-measures experimental design in which the same participants listen to all five stimuli; and (ii) functional localization at the single participant level, which establishes functional rather than anatomical correspondence across participants and therefore confers higher sensitivity (Nieto-Castañón and Fedorenko, 2012).

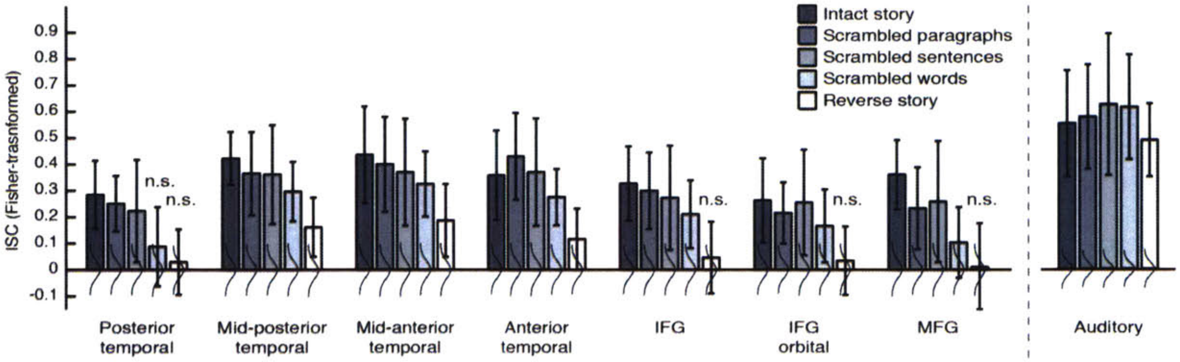

Figure 4.2 Inter-subject correlations in high-level language fROIs (left) and an auditory region (right), for stimuli with varying levels of scrambling. Each bar group shows Fisher-transformed ISCs in one region measured for 5 experimental conditions. Bars are resented in order of increasing stimulus scrambling (from dark to bright). Error bars show standard deviations across participants. Vertical thin curves show Gaussian fits to empirical null distributions against which each bar can be tested (n.s., non-significant results at a threshold of 0.05; FDR-corrected).

4.3.2. A less lenient approach reveals a different temporal integration hierarchy

In the previous section, inferences regarding regional integration time-windows were based on the observations that ISCs for one condition were significant whereas those for another condition were not. However, such tests do not license the conclusion that ISCs for these conditions are significantly different form one another (Nieuwenhuis et al., 2011). Conversely, the observation that ISCs for two different conditions are each significant does not license the conclusion that those ISCs are indistinguishable. Therefore, we next tested the ISCs for each scrambled stimulus directly against ISCs for the intact story. We asked what level of stimulus scrambling was sufficient to decrease input tracking below the level of tracking observed for fully intact linguistic input.

We found that only reverse-story ISCs were weaker than intact-story ISCs in the mid-anterior temporal ($t_{(16)}$=5.23, Cohen's d=1.31, p=0.002), anterior temporal ($t_{(16)}$=5.12, d=1.28, p=0.003), and the orbital IFG ($t_{(16)}$=3.89, d=0.97, p=0.009) fROIs. Based on these tests, activity in these fROIs exhibits highly reliable input tracking as long as stimuli are well structured at the scale of morphemes or single words and, hence, that is their integration time-window. Note that this temporal scale is larger than the one suggested by the previous analysis, which focused on the significance of the reverse-story ISCs.

In three other fROIs, a coarser level of stimulus scrambling was already sufficient to decrease input tracking, with scrambled-words ISCs being weaker than intact-story ISCs. These fROIs were located in the posterior temporal cortex ($t_{(16)}$=4.29, d=1.07, p=0.005), mid-posterior temporal cortex ($t_{(16)}$=4.44, d=1.11, p=9x10^{-4}), and IFG ($t_{(16)}$=3.35, d=0.84, p=0.016). By the standards of the current analysis, activity in these fROIs exhibits reliable input tracking as long as stimuli are well structured at the scale of sentences and, hence, that is their integration time-window. For the two latter fROIs, this time-window is again longer than that suggested by the previous analysis.

Finally, in the MFG fROI, scrambled paragraphs were already ill-structured enough to decrease ISCs compared to the intact story ($t_{(16)}$=3.11, d=0.78, p=0.032). Nevertheless, reordered sentences, which are even more incoherent than reordered paragraphs, did not produce ISCs that were significantly weaker than intact-story ISCs ($t_{(16)}$=1.55, d=0.39, p=0.46). Therefore, this analysis cannot conclusively characterize the integration time-window of this fROI.

4.3.3. Would a more rigorous approach have enough statistical power?

The previous analysis directly compared ISCs across different conditions, and concluded that different fROIs had different integration time-windows. However, such inferences are only licensed by direct tests of region-by-condition interactions (Nieuwenhuis et al., 2011). Namely, the finding that a difference between ISCs for two conditions is significant in one fROI but not in another does not itself imply that these two fROIs differ in their pattern of ISCs. Therefore, our main analysis aimed to explicitly compare how ISCs depended on stimulus scrambling across fROIs.

To demonstrate that such an approach will have enough power to detect functional differences across regions, we first applied it to a dataset for which we hypothesized that such differences must exist. Namely, we compared ISC patterns between our language fROIs and early auditory areas: whereas the former track the intact story more closely than the reverse story, the latter should similarly track both stimuli. To this end, we computed auditory ISCs based on BOLD time-series from an anatomically defined area covering Heschl's gyrus in the left hemisphere. Consistent with our hypothesis, a two-way, repeated-measures ANOVA with region (7 language fROIs + 1 auditory region) and condition (5 levels) as within-participant factors revealed a significant interaction ($F_{(28,448)}$=2.51, partial η^2=0.14, p=10^{-4}). To further explore this interaction, we performed a series of 2x2 ANOVAs that each compared intact-story ISCs and reverse-story ISCs between the auditory region and each language fROI. We found that six language regions showed stronger differences between story-intact ISCs and story-reverse ISCs compared to the auditory region (for all tests, $F>9.04$, partial $\eta^2>0.35$, $p<0.043$). Only the orbital IFG fROI did not differ from the auditory regions. Therefore, in our sample, a direct test of a region-by-condition interaction has sufficient power.

4.3.4. A region-by-condition interaction test finds no temporal integration hierarchy

When comparing functional responses across brain regions, it is critical to take into account regional differences in baseline responsiveness, because these might explain away region-by-stimulus interactions (e.g., Blank et al., 2016). Namely, if one region tracks fully coherent inputs more strongly than another region, then it has a bigger dynamic range for modulating input tracking and might exhibit seemingly larger decrements in tracking for scrambled stimuli. Thus, whereas an ANOVA might conclude that a change from an intact-story ISC of 0.5 to a reverse-story ISC of 0.25 is stronger than a change from 0.2 to 0.1, they both constitute a 50% decrease. Conversely, not accounting for baseline differences across regions might mask region-by-condition interactions when they do exist. Thus, whereas an ANOVA might conclude that a change from an intact-story ISC of 0.5 to a reverse-story ISC of 0.4 is similar to a change from 0.2 to 0.1, the former constitutes only a 20% decrease.

Indeed, the language fROIs differ from one another in their "baseline" input tracking: a one-way ANOVA performed on the intact-story ISCs with fROI (7 levels) as a within-participant factor revealed a significant main effect ($F_{(6,96)}=4.36$, partial $\eta^2=0.21$, $p=5\times10^{-4}$). Specifically, the mid-posterior temporal fROI exhibited stronger tracking than the posterior temporal ($t_{(16)}=3.86$, $d=0.96$, $p=0.03$), IFG ($t_{(16)}=4.01$, $d=1.00$, $p=0.024$), and orbital IFG ($t_{(16)}=4.44$, $d=1.11$, $p=0.029$) fROIs. The mid-anterior fROI also exhibited stronger tracking than the orbital IFG fROI ($t_{(16)}=4.40$, $d=1.10$, $p=0.037$). Therefore, prior to testing for region-by-condition interactions, we needed to control for such baseline regional differences.

To this end, we first regressed regional ISCs for all four scrambled conditions against regional ISCs for the intact-story in a linear, mixed-effects regression. The fixed effects in the model additionally included condition and an interaction between condition and story-intact ISCs. The random by-participant effects included an intercept and slopes for all fixed effects. Next, fitted estimates of regional ISCs for scrambled stimuli were computed using only the coefficients for story-intact ISCs and their interactions (i.e., ignoring the coefficients for condition). Finally, these fitted estimates were subtracted from their respective ISCs to obtain residuals that were not contaminated by regional baseline differences.

These regional ISC-residuals for scrambled conditions, along with the regional intact-story ISCs, were then submitted to a two-way, repeated-measures ANOVA. This test showed only a trend towards a region-by-condition interaction ($F_{(24,384)}$=1.43, partial η^2=0.08, p=0.087). To explore this trend, we performed a series of two-way ANOVAs that included all fROIs but were each restricted to a specific pair of conditions, consisting of one scrambled stimulus and the intact story. No test revealed a region-by-condition interaction. However, prior to FDR-correction, there was a trend towards a regional difference in the effect of scrambled words ($F_{(6,96)}$=2.17, partial η^2=0.12, p=0.053). A series of 2×2 ANOVAs comparing this effect across pairs of fROIs again revealed no significant results. Prior to FDR-correction, the only significant differences suggested that the MFG showed a stronger decrease in ISCs to scrambled words compared to the mid-posterior temporal ($F_{(1,16)}$=7.04, partial η^2=0.31, p=0.012), anterior temporal ($F_{(6,96)}$=7.04, partial η^2=0.3, p=0.019), IFG ($F_{(6,96)}$=10.78, partial η^2=0.4, p=0.006), and orbital IFG ($F_{(6,96)}$=7.75, partial η^2=0.33, p=0.016) fROIs. We note that another region-by-condition interaction, significant prior to FDR correction, was observed for the difference between intact-story ISCs and scrambled-paragraphs ISCs. However, this interaction was caused by one fROI showing a descriptive increase in ISCs for scrambled paragraphs (Figure 4.2) – a pattern that, we believe, reflects noise in the data and cannot be otherwise parsimoniously explained.

To relax our criteria for a region-by-condition interaction, we performed these ANOVAs again on the raw ISCs without controlling for regional differences in baseline input tracking. Now, the interaction reached significance ($F_{(24,384)}$=1.61, partial η^2=0.09, p=0.035). This interaction was driven by a differential effect of the scrambled words condition relative to the intact story across fROIs ($F_{(6,96)}$=3.78, partial η^2=0.19, p=0.017). As before, the MFG showed a stronger decrease in ISCs to scrambled words compared to the mid-posterior temporal ($F_{(1,16)}$=18.84, partial η^2=0.54, p=0.039), anterior temporal ($F_{(6,96)}$=14.8, partial η^2=0.48, p=0.036), and IFG ($F_{(6,96)}$=16.07, partial η^2=0.5, p=0.039) fROIs. We note that when we repeated the full ANOVA while excluding the MFG, the region-by-condition interaction was no longer significant ($F_{(20,320)}$=1.22, partial η^2=0.07, p=0.234).

In conclusion, an explicit test for a region-by-condition interaction does not provide strong evidence for a temporal integration hierarchy across language fROIs. With relaxed statistical standards, the only functional difference across regions distinguishes the MFG from the rest of the network, because its input tracking is more adversely affected by scrambling words, i.e., by the lack of phrase and sentence structures. In contrast, the four fROIs in the temporal cortex and the two inferior frontal regions are all affected indistinguishably by input scrambling and therefore show a common profile of temporal integration. To characterize this profile, we explored the main effect of condition in our main analysis ($F_{(4,64)}$=16.5, partial η^2=0.51, p=2×10^{-9}). A series of ANOVAs comparing each pair of conditions across all fROIs revealed that ISCs to the intact story and scrambled paragraphs did not differ ($F_{(1,16)}$=1.73, partial η^2=0.1, p=0.76), but were stronger than ISCs to scrambled words (for both tests, $F_{(1,16)}$>8.91, partial η^2>0.36, p<8×10^{-4}), with ISCs to the scrambled sentences being intermediate between these conditions (for all tests, $F_{(1,16)}$<3.77, partial η^2<0.19, p>0.29). ISCs for the reverse story were the lowest (for all tests, $F_{(1,16)}$>20.38, partial η^2>0.56, p<0.002).

4.4. Discussion

The current study examined how reliably regions in the language network track linguistic inputs that are scrambled at varying levels. We recorded regional activity time-courses elicited by increasingly scrambled variants of a narrated story, measured their reliability across individuals, and found that left temporal and inferior frontal language regions all exhibited indistinguishable patterns of sensitivity to linguistic structure at different grains. Namely, these regions all tracked scrambled paragraphs as reliably as they tracked the intact story, but tracked scrambled words less reliably, with intermediately reliable tracking of scrambled sentences. Therefore, these regions integrate information over a common temporal window that is not sensitive to structure across paragraphs but is sensitive to structure across words. This integration time-window reflects sentence-level processes and, perhaps, discourse processes that operate across small neighborhoods of sentences. Moreover, language regions tracked scrambled words more reliably than a reverse story, so their integration time-window also indicates the involvement of lexical and/or morphological processes. This common profile of temporal integration provides a novel functional signature of perisylvian, high-level language regions.

Given that we demonstrated sufficient statistical power to detect potential differences across regions in their functional profiles, we suggest that the absence of evidence for a functional dissociation across language regions should be taken as evidence for absence. Our results are inconsistent with a division of linguistic labor that is topographically organized by integration timescales across the language network (Dehaene-Lambertz et al., 2006; Stephens et al., 2013). Instead, this network plays a unique functional role within a broader cortical hierarchy for temporal integration, following lower-level speech-processing regions (Poeppel, 2003; Vagharchakian et al., 2012; Mesgarani et al., 2014; Overath et al., 2015) and preceding higher-level regions that integrate full narratives (Honey et al., 2012; Stephens et al., 2013; Chen et al., 2015). The only region of the language network that might occupy a different stage in a hierarchy of temporal integration is the MFG, exhibiting a potentially longer integration time-window.

4.4.1. Evidence for a distributed cognitive architecture of language processing

The finding that language regions have a common profile of temporal integration constrains cognitive models of language mechanisms, arguing against a functional dissociation between processes operating on linguistic representations of different sizes. Instead, it suggests that different regions exhibit shared, multiple timescales for linguistic processing, from sub-lexical to sentence-level integration. This notion is in agreement with linguistic and psycholinguistic theories suggesting language processing operates over a representational continuum from phonemes and morphemes through words with their syntactic and semantic attributes to phrase structures and their meanings (Joshi et al., 1975; Bresnan, 1982; Schabes et al., 1988; Pollard and Sag, 1994; Chomsky, 1995; Goldberg, 1995; Bybee, 1998, 2010; Jackendoff, 2002, 2007; Culicover and Jackendoff, 2005). It is also consistent with findings that the language processor keeps track of information that straddles past distinctions between language sounds, words and grammar (Clifton et al., 1984; MacDonald et al., 1994; Trueswell et al., 1994; Garnsey et al., 1997; Traxler et al., 2002; Farmer et al., 2006; Reali and Christiansen, 2007; Gennari and MacDonald, 2008; Schmidtke et al., 2014), and that even fine-grain representations are sensitive to broad linguistic contexts (Bradlow and Bent, 2008; Maye et al., 2008; Trude and Brown-Schmidt, 2012).

Even if different language regions all show the same profile of temporal integration, each region might implement a distinct set of processes relevant to such integration. Nonetheless, our finding of a functional characteristic that is distributed across the language network adds to prior neuroimaging studies reporting distributed activations, overlapping across diverse linguistic tasks (Gernsbacher and Kaschak, 2003; Démonet et al., 2005; Vigneau et al., 2006; Price, 2012). These include tasks of phonological (Scott and Wise, 2004; Hickok and Poeppel, 2007; Turkeltaub and Coslett, 2010), lexical (Paulesu et al., 1993; Indefrey and Levelt, 2004; Blumstein, 2009), syntactic (Caplan, 2007; Bautista and Wilson, 2016; Blank et al., 2016) and semantic (Bookheimer, 2002; Thompson-Schill, 2003; Patterson et al., 2007; Binder et al., 2009) processing. Unlike these traditional task-based studies, which used carefully controlled manipulations contrived to isolate particular aspects of linguistic processing, the current study employed an alternative approach (Hasson et al., 2004) based on richly structured stimuli in a naturalistic listening paradigm. It therefore importantly complements the prior evidence for a distributed architecture of language processing within which the very same neural circuits might support several different kinds of computations.

4.4.2. Why functional dissociations within the language network might go undetected

While we interpret our findings as supporting distributed processing throughout the language network, they are not inconsistent with functional dissociations across regions of this network. Indeed, neuropsychological findings indicate, despite their inconsistencies, that at least some language regions are specialized for some linguistic processes and not others, because some patients drastically differ from one another in their deficits (Caramazza and Coltheart, 2006). Therefore, it is possible that some fMRI evidence for distributed linguistic processing underestimate a more complex functional architecture within the language network.

For instance, a regional profile of temporal integration might be the sum of several different profiles originating from small, functionally distinct, neighboring subdivisions within it that cannot be well differentiated due to limited spatial resolution. However, given that the neuropsychological literature has studied many patients with lesions much larger than the spatial grain of fMRI, at least some functional subdivisions within the language network should be detectable with current neuroimaging methods. An alternative account for a common profile of temporal integration might be the rapid spread of linguistic information across distinct functional regions. Such transfer is likely, given the anatomical connections and functional synchronization among language regions (Saur et al., 2008; Blank et al., 2014). Thus, by the time the temporally slow fMRI BOLD signal is detected, an initially focal neural response might already appear ubiquitous throughout the language network.

Beyond the methodological limitations of fMRI, studying the division of linguistic labor across the language network also faces theoretical challenges. As mentioned above, traditional distinctions between psycholinguistic constructs (e.g., "words" vs. "rules") are no longer advocated by contemporary theories – yet they still guide neuroimaging studies (e.g., Friederici, 2012). Moreover, it remains unclear what functional distinctions should instead be tested. For instance, despite strikingly different linguistic deficits across some patients in the neuropsychological literature, the precise nature of these deficits is still under debate (Caplan et al., 1996, 2007, 2013; Dronkers, 2000; Caramazza et al., 2001; Wilson and Saygın, 2004; Hillis, 2007; Grodzinsky and Santi, 2008). Perhaps, then, neuroimaging investigations of the division of linguistic labor could better constrain cognitive models when their theoretical motivation is more grounded in such models.

4.4.3. A key methodology for neuroimaging studies of language processing

The current study tested whether the previously reported cortical hierarchy of integration time-windows (Lerner et al., 2011) corresponded to the language network, and demonstrated that language regions all play a similar functional role along this hierarchy. Why, then, did the original topography of this hierarchy appear to overlap with the language network?

We believe that this overlap was an illusory effect due, in large part, to relying on anatomical criteria for mapping the temporal integration hierarchy. Namely, this hierarchy was characterized by computing ISCs on a voxel-by-voxel basis, assuming that a given voxel (i.e., anatomical location) corresponded to a single functional unit across individuals. This assumption has been shown to be invalid in the temporal and frontal lobes, where high inter-individual variability in functional-to-anatomical mapping is evident (Amunts et al., 1999; Fischl et al., 2008; Tahmasebi et al., 2012) and language regions lie adjacent to other functional regions (Wise et al., 2001; Chein et al., 2002; Fedorenko et al., 2012a). Therefore, in areas where language regions happened to overlap across participants, high ISCs would be found even for relatively scrambled stimuli. In contrast, in areas that were more functionally heterogeneous across participants lower ISCs would be found, especially for scrambled stimuli that are tracked with different levels of reliability by networks with different integration timescales. These effects would seem to falsely indicate that different language regions occupy distinct positions in a hierarchy of temporal integration.

Here, we took inter-individual variability in the precise anatomical location of language regions into account by localizing these regions separately in each participant, and functional differences across language regions consequently dissolved. This methodology, together with the use of rigorous statistical tests for detecting functional dissociations across regions, was fundamental to establishing a profile of temporal integration as a functional signature of the language network. This study thus joins an emerging literature that employs powerful localization methods to characterize the functional architecture of the language network, calling many traditional views into question (Fedorenko et al., 2011; Fedorenko et al., 2012a; Blank et al., 2014; Fedorenko and Thompson-Schill, 2014; Blank et al., 2016).

4.4.4. Conclusion

As linguistic inputs unfold over time, we integrate them into structured representations that mediate language comprehension. Whereas such integration might proceed hierarchically across the cortex (at least for auditory processing if not for language in particular), high-level language regions in the temporal and inferior frontal cortex all exhibit a similar functional profile within this hierarchy. This finding suggests that the division of linguistic labor across the language network is not topographically organized according to distinct timescales of temporal integration or, by extension, distinct grains of linguistic representation. Instead, a distributed set of regions might all integrate linguistic information over common, multiple timescales.

4.5. Acknowledgements

We thank Uri Hasson (Psychology Department and the Princeton Neuroscience Institute, Princeton University) for providing the stimuli. We also thank the Athinoula A. Martinos Imaging Center at McGovern Institute for Brain Research at MIT, and the support team (Steve Shannon, Atsushi Takahashi, and Sheeba Arnold). E.F. was supported by a K99/R00 award HD 057522 from NICHD.

References

Amunts K, Schleicher A, Bürgel U, Mohlberg H, Uylings H, Zilles K (1999) Broca's region revisited: Cytoarchitecture and intersubject variability. Journal of Comparative Neurology 412:319-341.

Anderson M, Braak CT (2003) Permutation tests for multi-factorial analysis of variance. Journal of Statistical Computation and Simulation 73:85-113.

Barr DJ, Levy R, Scheepers C, Tily HJ (2013) Random effects structure for confirmatory hypothesis testing: Keep it maximal. Journal of Memory and Language 68:255-278.

Bates E, Wilson SM, Saygin AP, Dick F, Sereno MI, Knight RT, Dronkers NF (2003) Voxel-based lesion-symptom mapping. Nature Neuroscience 6:448-450.

Bautista A, Wilson SM (2016) Neural responses to grammatically and lexically degraded speech. Language, Cognition and Neuroscience 31:567-574.

Benjamini Y, Yekutieli D (2001) The control of the false discovery rate in multiple testing under dependency. Annals of statistics:1165-1188.

Binder JR, Desai RH, Graves WW, Conant LL (2009) Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. Cerebral Cortex 19:2767-2796.

Binder JR, Frost JA, Hammeke TA, Bellgowan PSF, Springer JA, Kaufman JN, Possing ET (2000) Human temporal lobe activation by speech and nonspeech sounds. Cerebral Cortex 10:512-528.

Blank I, Kanwisher N, Fedorenko E (2014) A functional dissociation between language and multiple-demand systems revealed in patterns of BOLD signal fluctuations. Journal of Neurophysiology 112:1105-1118.

Blank I, Balewski Z, Mahowald K, Fedorenko E (2016) Syntactic processing is distributed across the language system. NeuroImage 127:307-323.

Blumstein SE (2009) Auditory word recognition: Evidence from aphasia and functional neuroimaging. Language and Linguistics Compass 3:824-838.

Bookheimer S (2002) Functional MRI of language: New approaches to understanding the cortical organization of semantic processing. Annual Review of Neuroscience 25:151-188.

Bradlow AR, Bent T (2008) Perceptual adaptation to non-native speech. Cognition 106:707-729.

Braze D, Mencl WE, Tabor W, Pugh KR, Constable RT, Fulbright RK, Magnuson JS, Van Dyke JA, Shankweiler DP (2011) Unification of sentence processing via ear and eye: An fMRI study. Cortex 47:416-431.

Bresnan J (1982) The Mental Representation of Grammatical Relations. Cambridge, MA: MIT Press.

Bybee J (1998) A functionalist approach to grammar and its evolution. Evolution of Communication 2:249-278.

Bybee J (2010) Language, Usage and Cognition. Cambridge, UK: Cambridge University Press.

Caplan D (2007) Functional neuroimaging studies of syntactic processing in sentence comprehension: A critical selective review. Language and Linguistics Compass 1:32-47.

Caplan D, Hildebrandt N, Makris N (1996) Location of lesions in stroke patients with deficits in syntactic processing in sentence comprehension. Brain 119:933-949.

Caplan D, Michaud J, Hufford R (2013) Dissociations and associations of performance in syntactic comprehension in aphasia and their implications for the nature of aphasic deficits. Brain and Language 127:21-33.

Caplan D, Waters G, DeDe G, Michaud J, Reddy A (2007) A study of syntactic processing in aphasia I: Behavioral (psycholinguistic) aspects. Brain and Language 101:103-150.

Caramazza A, Coltheart M (2006) Cognitive neuropsychology twenty years on. Cognitive Neuropsychology 23:3-12.

Caramazza A, Capitani E, Rey A, Berndt RS (2001) Agrammatic Broca's aphasia is not associated with a single pattern of comprehension performance. Brain and Language 76:158-184.

Chein JM, Fissell K, Jacobs S, Fiez JA (2002) Functional heterogeneity within Broca's area during verbal working memory. Physiology and Behavior 77:635-639.

Chen J, Honey CJ, Simony E, Arcaro MJ, Norman KA, Hasson U (2015) Accessing real-life episodic information from minutes versus hours earlier modulates hippocampal and high-order cortical dynamics. Cerebral Cortex:bhv155.

Chomsky N (1995) The Minimalist Program. Cambridge, MA: MIT Press.

Clifton C, Frazier L, Connine C (1984) Lexical expectations in sentence comprehension. Journal of Verbal Learning and Verbal Behavior 23:696-708.

Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, Quigley MA, Meyerand ME (2001) Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. American Journal of Neuroradiology 22:1326-1333.

Culicover PW, Jackendoff R (2005) Simpler Syntax. Oxford, UK: Oxford University Press.

Dapretto M, Bookheimer SY (1999) Form and content: Dissociating syntax and semantics in sentence comprehension. Neuron 24:427-432.

Dehaene - Lambertz G, Dehaene S, Anton JL, Campagne A, Ciuciu P, Dehaene GP, Denghien I, Jobert A, LeBihan D, Sigman M (2006) Functional segregation of cortical language areas by sentence repetition. Human Brain Mapping 27:360-371.

Démonet J-F, Thierry G, Cardebat D (2005) Renewal of the neurophysiology of language: Functional neuroimaging. Physiological Reviews 85:49-95.

Dronkers NF (2000) The gratuitous relationship between Broca's aphasia and Broca's area. Behavioral and Brain Sciences 23:30-31.

Farmer TA, Christiansen MH, Monaghan P (2006) Phonological typicality influences on-line sentence comprehension. Proceedings of the National Academy of Sciences 103:12203-12208.

Fedorenko E (2014) The role of domain-general cognitive control in language comprehension. Frontiers in Psychology 5:335.

Fedorenko E, Thompson-Schill SL (2014) Reworking the language network. Trends in Cognitive Sciences 18:120-126.

Fedorenko E, Varley R (2016) Language and thought are not the same thing: Evidence from neuroimaging and neurological patients. Annals of the New York Academy of Sciences 1369:132-153.

Fedorenko E, Behr MK, Kanwisher N (2011) Functional specificity for high-level linguistic processing in the human brain. Proceedings of the National Academy of Sciences 108:16428-16433.

Fedorenko E, Duncan J, Kanwisher N (2012a) Language-selective and domain-general regions lie side by side within Broca's area. Current Biology 22:2059-2062.

Fedorenko E, Nieto-Castañón A, Kanwisher N (2012b) Lexical and syntactic representations in the brain: an fMRI investigation with multi-voxel pattern analyses. Neuropsychologia 50:499-513.

Fedorenko E, Hsieh P-J, Balewski Z (2015) A possible functional localiser for identifying brain regions sensitive to sentence-level prosody. Language, Cognition and Neuroscience 30:120-148.

Fedorenko E, Hsieh P-J, Nieto-Castañón A, Whitfield-Gabrieli S, Kanwisher N (2010) New method for fMRI investigations of language: Defining ROIs functionally in individual subjects. Journal of Neurophysiology 104:1177-1194.

Fischl B, Rajendran N, Busa E, Augustinack J, Hinds O, Yeo BTT, Mohlberg H, Amunts K, Zilles K (2008) Cortical folding patterns and predicting cytoarchitecture. Cerebral Cortex 18:1973-1980.

Friederici AD (2002) Towards a neural basis of auditory sentence processing. Trends in Cognitive Sciences 6:78-84.

Friederici AD (2012) The cortical language circuit: From auditory perception to sentence comprehension. Trends in Cognitive Sciences 16:262-268.

Garnsey SM, Pearlmutter NJ, Myers E, Lotocky MA (1997) The contributions of verb bias and plausibility to the comprehension of temporarily ambiguous sentences. Journal of Memory and Language 37:58-93.

Gennari SP, MacDonald MC (2008) Semantic indeterminacy in object relative clauses. Journal of Memory and Language 58:161-187.

Gernsbacher MA, Kaschak MP (2003) Neuroimaging studies of language production and comprehension. Annual Review of Psychology 54:91-114.

Gill PMW (2007) Efficient calculation of p-values in linear-statistic permutation significance tests. Journal of Statistical Computation and Simulation 77:55-61.

Gloor P (1997) The temporal lobe and limbic system. New York, NY: Oxford University Press.

Goldberg AE (1995) Constructions: A construction grammar approach to argument structure. Chicago, IL: University of Chicago Press.

Golland Y, Bentin S, Gelbard H, Benjamini Y, Heller R, Nir Y, Hasson U, Malach R (2007) Extrinsic and intrinsic systems in the posterior cortex of the human brain revealed during natural sensory stimulation. Cerebral Cortex 17:766-777.

Grodzinsky Y, Santi A (2008) The battle for Broca's region. Trends in Cognitive Sciences 12:474-480.

Hagoort P (2013) MUC (memory, unification, control) and beyond. Frontiers in Psychology 4:416.

Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R (2004) Intersubject synchronization of cortical activity during natural vision. science 303:1634-1640.

Hasson U, Avidan G, Gelbard H, Vallines I, Harel M, Minshew N, Behrmann M (2009) Shared and idiosyncratic cortical activation patterns in autism revealed under continuous real - life viewing conditions. Autism Research 2:220-231.

Hickok G, Poeppel D (2004) Dorsal and ventral streams: A framework for understanding aspects of the functional anatomy of language. Cognition 92:67-99.

Hickok G, Poeppel D (2007) The cortical organization of speech processing. Nature Reviews Neuroscience 8:393-402.

Hillis AE (2007) Aphasia progress in the last quarter of a century. Neurology 69:200-213.

Honey CJ, Thesen T, Donner TH, Silbert LJ, Carlson CE, Devinsky O, Doyle WK, Rubin N, Heeger DJ, Hasson U (2012) Slow cortical dynamics and the accumulation of information over long timescales. Neuron 76:423-434.

Humphries C, Binder JR, Medler DA, Liebenthal E (2006) Syntactic and semantic modulation of neural activity during auditory sentence comprehension. Journal of Cognitive Neuroscience 18:665-679.

Indefrey P, Levelt WJM (2004) The spatial and temporal signatures of word production components. Cognition 92:101-144.

Jackendoff R (2002) Foundation of Language: Brain, Meaning, Grammar, Evolution. Oxford, UK: Oxford University Press.

Jackendoff R (2007) A parallel architecture perspective on language processing. Brain Research 1146:2-22.

Joshi AK, Levy LS, Takahashi M (1975) Tree adjunct grammars. Journal of Computer and System Sciences 10:136-163.

Keller TA, Carpenter PA, Just MA (2001) The neural bases of sentence comprehension: a fMRI examination of syntactic and lexical processing. Cerebral Cortex 11:223-237.

Kimura D, Folb S (1968) Neural processing of backwards-speech sounds. Science 161:395-396.

Koeda M, Takahashi H, Yahata N, Matsuura M, Asai K, Okubo Y, Tanaka H (2006) Language processing and human voice perception in schizophrenia: A functional magnetic resonance imaging study. Biological Psychiatry 59:948-957.

Lerner Y, Honey CJ, Silbert LJ, Hasson U (2011) Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. The Journal of Neuroscience 31:2906-2915.

MacDonald MC, Pearlmutter NJ, Seidenberg MS (1994) The lexical nature of syntactic ambiguity resolution. Psychological Review 101:676-703.

Maye J, Aslin RN, Tanenhaus MK (2008) The weckud wetch of the wast: Lexical adaptation to a novel accent. Cognitive Science 32:543-562.

Mazoyer BM, Tzourio N, Frak V, Syrota A, Murayama N, Levrier O, Salamon G, Dehaene S, Cohen L, Mehler J (1993) The cortical representation of speech. Journal of Cognitive Neuroscience 5:467-479.

Menenti L, Gierhan SME, Segaert K, Hagoort P (2011) Shared language overlap and segregation of the neuronal infrastructure for speaking and listening revealed by functional MRI. Psychological Science 22:1173-1182.

Mesgarani N, Cheung C, Johnson K, Chang EF (2014) Phonetic feature encoding in human superior temporal gyrus. Science 343:1006-1010.

Mesulam MM, Thompson CK, Weintraub S, Rogalski EJ (2015) The Wernicke conundrum and the anatomy of language comprehension in primary progressive aphasia. Brain 138:2423-2437.

Mewhort DJK, Johns BT, Kelly M (2010) Applying the permutation test to factorial designs. Behavior Research Methods 42:366-372.

Mirman D, Chen Q, Zhang Y, Wang Z, Faseyitan OK, Coslett HB, Schwartz MF (2015) Neural organization of spoken language revealed by lesion-symptom mapping. Nature Communications 6.

Nieto-Castañón A, Fedorenko E (2012) Subject-specific functional localizers increase sensitivity and functional resolution of multi-subject analyses. NeuroImage 63:1646-1669.

Nieuwenhuis S, Forstmann BU, Wagenmakers E-J (2011) Erroneous analyses of interactions in neuroscience: A problem of significance. Nature Neuroscience 14:1105-1107.

Norman-Haignere S, Kanwisher N, McDermott JH (2015) Distinct cortical pathways for music and speech revealed by hypothesis-free voxel decomposition. Neuron 88:1281-1296.

Overath T, McDermott JH, Zarate JM, Poeppel D (2015) The cortical analysis of speech-specific temporal structure revealed by responses to sound quilts. Nature Neuroscience 18:903-911.

Patterson K, Nestor PJ, Rogers TT (2007) Where do you know what you know? The representation of semantic knowledge in the human brain. Nature Reviews Neuroscience 8:976-987.

Paulesu E, Frith CD, Frackowiak RSJ (1993) The neural correlates of the verbal component of working memory. Nature 362:342-345.

Poeppel D (2003) The analysis of speech in different temporal integration windows: Cerebral lateralization as 'asymmetric sampling in time'. Speech Communication 41:245-255.

Poldrack RA (2006) Can cognitive processes be inferred from neuroimaging data? Trends in Cognitive Sciences 10:59-63.

Pollard C, Sag IA (1994) Head-driven phrase structure grammar. Chicago, IL: University of Chicago Press.

Price CJ (2012) A review and synthesis of the first 20 years of PET and fMRI studies of heard speech, spoken language and reading. NeuroImage 62:816-847.

Reali F, Christiansen MH (2007) Processing of relative clauses is made easier by frequency of occurrence. Journal of Memory and Language 57:1-23.

Saur D, Kreher BW, Schnell S, Kümmerer D, Kellmeyer P, Vry M-S, Umarova R, Musso M, Glauche V, Abel S (2008) Ventral and dorsal pathways for language. Proceedings of the National Academy of Sciences 105:18035-18040.

Schabes Y, Abeille A, Joshi AK (1988) Parsing strategies with'lexicalized'grammars: application to tree adjoining grammars. In: Proceedings of the 12th conference on Computational linguistics-Volume 2, pp 578-583: Association for Computational Linguistics.

Schmidtke D, Conrad M, Jacobs AM (2014) Phonological iconicity. Frontiers in Psychology 5:80.

Scott SK, Wise RJS (2004) The functional neuroanatomy of prelexical processing in speech perception. Cognition 92:13-45.

Silver NC, Dunlap WP (1987) Averaging correlation coefficients: should Fisher's z transformation be used? Journal of Applied Psychology 72:146.

Stephens GJ, Honey CJ, Hasson U (2013) A place for time: The spatiotemporal structure of neural dynamics during natural audition. Journal of Neurophysiology 110:2019-2026.

Tahmasebi AM, Davis MH, Wild CJ, Rodd JM, Hakyemez H, Abolmaesumi P, Johnsrude IS (2012) Is the link between anatomical structure and function equally strong at all cognitive levels of processing? Cerebral Cortex 22:1593-1603.

Theiler J, Eubank S, Longtin A, Galdrikian B, Farmer JD (1992) Testing for nonlinearity in time series: The method of surrogate data. Physica D: Nonlinear Phenomena 58:77-94.

Thompson-Schill SL (2003) Neuroimaging studies of semantic memory: Inferring "how" from "where". Neuropsychologia 41:280-292.

Traxler MJ, Morris RK, Seely RE (2002) Processing subject and object relative clauses: Evidence from eye movements. Journal of Memory and Language 47:69-90.

Trude AM, Brown-Schmidt S (2012) Talker-specific perceptual adaptation during online speech perception. Language and Cognitive Processes 27:979-1001.

Trueswell JC, Tanenhaus MK, Garnsey SM (1994) Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. Journal of Memory and Language 33:285-318.

Turkeltaub PE, Coslett HB (2010) Localization of sublexical speech perception components. Brain and Language 114:1-15.

Ullman MT (2004) Contributions of memory circuits to language: The declarative/procedural model. Cognition 92:231-270.

Vagharchakian L, Dehaene-Lambertz G, Pallier C, Dehaene S (2012) A temporal bottleneck in the language comprehension network. The Journal of Neuroscience 32:9089-9102.

Vigneau M, Beaucousin V, Herve P-Y, Duffau H, Crivello F, Houde O, Mazoyer B, Tzourio-Mazoyer N (2006) Meta-analyzing left hemisphere language areas: Phonology, semantics, and sentence processing. NeuroImage 30:1414-1432.

Whitfield-Gabrieli S, Nieto-Castañón A (2012) Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. Brain connectivity 2:125-141.

Wilson SM, Saygın AP (2004) Grammaticality judgment in aphasia: Deficits are not specific to syntactic structures, aphasic syndromes, or lesion sites. Journal of Cognitive Neuroscience 16:238-252.

Wise RJS, Scott SK, Blank SC, Mummery CJ, Murphy K, Warburton EA (2001) Separate neural subsystems within "Wernicke's area". Brain 124:83-95.

[this page intentionally left blank]

Summary

In the work presented here, I characterized the functional architecture of language comprehension mechanisms from two, complementary perspectives. First, I critically compared the functional profiles of domain-specific language regions and domain-general MD regions, demonstrating that these regions form two functionally dissociated networks. Similar results in the neuropsychological literature (Broca, 1861/2006; Wernicke, 1874/1969; Luria et al., 1965; Varley and Siegal, 2000; Varley et al., 2005; Apperly et al., 2006; Bek et al., 2010; Willems et al., 2011) have established that these two networks can be independently impaired and, therefore, functional integration across them is not necessary; nonetheless, these results could not exclude the possibility that such integration is possible and common in the neurotypical brain. The current findings thus extend previous results in showing that even prior to any brain damage activity in the language network is not synchronized with activity in the MD network during naturalistic cognition. Language-specific and domain-general mechanisms likely support distinct computations, playing different roles in a functional ontology of language comprehension.

In addition, I provide a critical insight into the nature of these respective roles. Namely, I report that the language network closely tracks linguistic input with relatively little individual differences in its processing patterns whereas the MD network shows weak input tracking. The finding that MD network activity is not highly sensitive to features of the input during natural language processing challenges neuroimaging findings obtained with task-based designs (Stromswold et al., 1996; Stowe et al., 1998; Caplan et al., 1999; Fiez et al., 1999; Fiebach et al., 2002; Chee et al., 2003; Constable et al., 2004; Rodd et al., 2005; Chen et al., 2006; Nakic et al., 2006; Nieuwland et al., 2007; Novais-Santos et al., 2007; Hauk et al., 2008; Yarkoni et al., 2008; Carreiras et al., 2009; January et al., 2009; Peelle et al., 2009; Ye and Zhou, 2009; Barde et al., 2012; McMillan et al., 2012; McMillan et al., 2013) and importantly constrains the possible contributions of this network to comprehension. For instance, it might continuously allocate a pool of resources to support comprehension but modulate this allocation mostly according to internals factors affecting attention and focus, and less according to immediate external demands.

The second perspective for characterizing the functional architecture of comprehension mechanisms focused on the high-level language network, examining the division of linguistic labor across its constituent regions. Here, I revisited two central hypotheses regarding the organization of this network, claiming that in both cases neuroimaging evidence suggests distributed processing that (i) is spread throughout the network, and (ii) is overlapping across putative linguistic processes. In one study, I demonstrated that all language regions are indistinguishably sensitive to manipulations of syntactic complexity, which contradicts many previous neuroimaging studies (for a meta-analysis, see Friederici, 2011) but is consistent with neuropsychological findings (Caplan et al., 1996; Dick et al., 2001; Wilson and Saygin, 2004; Mesulam et al., 2015). In the other study, I found that language regions in the temporal and inferior frontal cortices all integrate linguistic information over a shared set of timescales, a result that is inconsistent with previous findings (Lerner et al., 2011). Thus, these regions all occupy a common position, rather than several distinct positions, along a cortical hierarchy of temporal integration. Both instances of distributed processing could be interpreted to suggest that multiple aspects of linguistic processing are cognitively inseparable from one another, in line with current psycholinguistic proposals (Joshi et al., 1975; Bresnan, 1982; Schabes et al., 1988; Pollard and Sag, 1994; Chomsky, 1995; Goldberg, 1995; Bybee, 1998, 2010; Jackendoff, 2002, 2007; Culicover and Jackendoff, 2005).

The observation that these findings challenge existing views in the neuroimaging literature but are nevertheless consistent with both neuropsychological findings and theoretical psycholinguistic advancements speaks to the methodological strengths of this work. As discussed in the Introduction, the current studies harness the complementary strengths of single-participant functional localization (Saxe et al., 2006; Fedorenko et al., 2010), data-driven methods using naturalistic paradigms (Biswal et al., 1995; Hasson et al., 2004), and rigorous statistical tests (Nieuwenhuis et al., 2011), thus providing novel means to answer old questions. In contrast, the methods used by many prior studies might have been insufficiently powerful to properly address these questions. Indeed, in some chapters I explicitly articulated how such methodological differences could have led to the fundamental discrepancies between the existing literature and the current work. Critically, the central methodological message of the dissertation holds even if one remains agnostic with regard to the theoretical implications of the reported findings: prior to interpreting any results, we need to trust that (i) group-level data accurately represents individual brains; (ii) measures are ecologically valid; and (iii) hypotheses are explicitly tested.

This point regarding "theoretical agnosticism" is worth emphasizing. Thus far, when discussing the results of this work, I have strongly advocated only those interpretations that required a relatively "simple" link from operational measures to the underlying psychological constructs (if such links in neuroimaging studies are ever simple). When proposing deeper interpretations that might be indirectly implied by the data, I have attempted to be careful and critical. Perhaps unexpectedly, I have resisted the free reign that I am allowed in the context of this dissertation and have not expressed my own thoughts on the functional ontology of language. The main reason for avoiding such musings is my attempt to let the data, rather than my intuitions, dictate my views. This inclination might appear trivial, or naïve; nonetheless, I believe that the modern history of linguistics has taught us that intuitions are both hard to ignore and dangerous, creating theoretical "attractor fields" that require great effort to destroy. Perhaps Wittgenstein (1953/2010) has already warned us against the challenge of intuitions when he claimed that we were trapped inside our language. Human minds, both the creators and investigators of language, are "rats who build the labyrinth from which they [...] try to escape" (Raymond Queneau, cited in Lescure, 1986).

Still, even with relatively minimal commitment to interpretation, I have demonstrated above that the current work importantly constrains the functional ontology of language comprehension mechanisms. In light of such insights, several questions warrant future investigation. First, what are the circumstances under which language and MD regions might interact, and what are the goals of such interactions? Second, what are the particular processes implemented in each network? (see, e.g., Fedorenko, 2014; Fedorenko and Thompson-Schill, 2014; one hypothesis has already been described here, in the context of syntactic processing). Third, what organizational principles govern the division of linguistic labor across language regions, given those aspects of comprehension that appear distributed throughout the network? Clearly, these questions reveal the few explicit intuitions I have regarding the functional architecture of comprehension mechanisms: that different networks engaged in comprehension should not operate in complete independence; and that multiple, anatomically distributed regions, even if internally synchronized, are not likely to be completely redundant in their respective computations.

Addressing these questions may require alternative methods to the ones used in the current work. For instance, if interactions between language and MD networks are brief, or if different language regions are active at different times during comprehension, then methods with a higher temporal resolution are of critical importance. Whereas some progress can be made by simply studying fMRI BOLD signal time-courses over shorter time bins (Chai et al., in press), qualitatively different methods that can measure neural activity with millisecond resolution (e.g., electrocorticoraphy; see Ojemann et al., 1989a; Ojemann et al., 1989b; Towle et al., 2008) are likely to be necessary. Ultimately, however, answers to these questions will depend on methodologies that not only offer high spatial and temporal resolution but also license causal inferences, such as transcranial magnetic and intra-operative stimulation techniques. Using these tools, some putative functional dissociations within the language network have already been suggested (Gough et al., 2005; Lee et al., in prep).

Beyond employing alternative methods, addressing the questions above might require cognitive neuroscientists of language to also replace the hypothesized psycholinguistic constructs that have been used to date. In other words, it is possible that the struggle to obtain reproducible functional dissociations across different regions that engage in comprehension has resulted from asking the wrong questions (Grodzinsky, 2010). Perhaps, then, gross functional distinctions such as the ones explored here should be refined into more subtle distinctions (e.g., Frankland and Greene, 2015). Alternatively, we might have to abandon the theoretical frameworks that have thus far guided neuroimaging studies of comprehension (e.g., Friederici, 2012) in favor of new frameworks informed by contemporary psycholinguistic proposals (e.g., Chomsky, 1995; Goldberg, 1995; Bybee, 1998; Jackendoff, 2002; Culicover and Jackendoff, 2005; Bybee, 2010). Unfortunately, such proposals often call for collapsing the functional distinctions that have been pursued by earlier theories and instead advocate for more unified functional architectures. Similarly, neuropsychological accounts of functional dissociations across patients with different linguistic deficits are yet to agree on a functional ontology of language that would explain these dissociations (Caplan et al., 1996, 2007, 2013; Dronkers, 2000; Caramazza et al., 2001; Wilson and Saygin, 2004; Hillis, 2007; Grodzinsky and Santi, 2008).

How are we to advance, then? One scientific strategy is to bring neuroimaging, psycholinguistics and neuropsychology closer together such that they frequently and deliberately inform each other rather than proceed in relative independence. For instance, data-driven neuroimaging methods could explore the sources of variability in neural responses to natural linguistic input, in order to recover the independent dimensions in such input that account for functional differences across regions engaged in comprehension (for a similar approach in auditory processing, see Norman-Haignere et al., 2015). Then, psycholinguists could attempt to interpret these dimensions and integrate them into coherent theories within their existing frameworks. Such interpretations could, in turn, be validated by further studies in all three disciplines. I believe this dissertation takes initial steps towards establishing such an enterprise of strong disciplinary inter-dependence in the search for a functional ontology of language.

References

Apperly IA, Samson D, Carroll N, Hussain S, Humphreys G (2006) Intact first-and second-order false belief reasoning in a patient with severely impaired grammar. Social Neuroscience 1:334-348.

Barde LH, Yeatman JD, Lee ES, Glover G, Feldman HM (2012) Differences in neural activation between preterm and full term born adolescents on a sentence comprehension task: Implications for educational accommodations. Developmental Cognitive Neuroscience 2:S114-S128.

Bek J, Blades M, Siegal M, Varley RA (2010) Language and spatial reorientation: Evidence from severe aphasia. Journal of Experimental Psychology: Learning, Memory, and Cognition 36:646-658.

Biswal B, Zerrin Yetkin F, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo - planar MRI. Magnetic Resonance in Medicine 34:537-541.

Bresnan J (1982) The Mental Representation of Grammatical Relations. Cambridge, MA: MIT Press.

Broca P (1861/2006) Comments regarding the seat of the faculty of spoken language, followed by an observation of aphemia (loss of speech). In: Broca's Region (Grodzinsky Y, Amunts K, eds), pp 291-304. New York, NY: Oxford University Press.

Bybee J (1998) A functionalist approach to grammar and its evolution. Evolution of Communication 2:249-278.

Bybee J (2010) Language, Usage and Cognition. Cambridge, UK: Cambridge University Press.

Caplan D, Hildebrandt N, Makris N (1996) Location of lesions in stroke patients with deficits in syntactic processing in sentence comprehension. Brain 119:933-949.

Caplan D, Alpert N, Waters G (1999) PET studies of syntactic processing with auditory sentence presentation. NeuroImage 9:343-351.

Caplan D, Michaud J, Hufford R (2013) Dissociations and associations of performance in syntactic comprehension in aphasia and their implications for the nature of aphasic deficits. Brain and Language 127:21-33.

Caplan D, Waters G, DeDe G, Michaud J, Reddy A (2007) A study of syntactic processing in aphasia I: Behavioral (psycholinguistic) aspects. Brain and Language 101:103-150.

Caramazza A, Capitani E, Rey A, Berndt RS (2001) Agrammatic Broca's aphasia is not associated with a single pattern of comprehension performance. Brain and Language 76:158-184.

Carreiras M, Riba J, Vergara M, Heldmann M, Münte TF (2009) Syllable congruency and word frequency effects on brain activation. Human Brain Mapping 30:3079-3088.

Chai LR, Mattar MG, Blank I, Fedorenko E, Bassett DS (in press) Functional network dynamics of the language system. Cerebral Cortex.

Chee MWL, Westphal C, Goh J, Graham S, Song AW (2003) Word frequency and subsequent memory effects studied using event-related fMRI. NeuroImage 20:1042-1051.

Chen E, West WC, Waters G, Caplan D (2006) Determinants of BOLD signal correlates of processing object-extracted relative clauses. Cortex 42:591-604.

Chomsky N (1995) The Minimalist Program. Cambridge, MA: MIT Press.

Constable RT, Pugh KR, Berroya E, Mencl WE, Westerveld M, Ni W, Shankweiler D (2004) Sentence complexity and input modality effects in sentence comprehension: an fMRI study. NeuroImage 22:11-21.

Culicover PW, Jackendoff R (2005) Simpler Syntax. Oxford, UK: Oxford University Press.

Dick F, Bates E, Wulfeck B, Utman JA, Dronkers N, Gernsbacher MA (2001) Language deficits, localization, and grammar: Evidence for a distributive model of language breakdown in aphasic patients and neurologically intact individuals. Psychological Review 108:759-788.

Dronkers NF (2000) The gratuitous relationship between Broca's aphasia and Broca's area. Behavioral and Brain Sciences 23:30-31.

Fedorenko E (2014) The role of domain-general cognitive control in language comprehension. Frontiers in Psychology 5:335.

Fedorenko E, Thompson-Schill SL (2014) Reworking the language network. Trends in Cognitive Sciences 18:120-126.

Fedorenko E, Hsieh P-J, Nieto-Castañón A, Whitfield-Gabrieli S, Kanwisher N (2010) New method for fMRI investigations of language: Defining ROIs functionally in individual subjects. Journal of Neurophysiology 104:1177-1194.

Fiebach CJ, Friederici AD, Müller K, Von Cramon DY (2002) fMRI evidence for dual routes to the mental lexicon in visual word recognition. Journal of Cognitive Neuroscience 14:11-23.

Fiez JA, Balota DA, Raichle ME, Petersen SE (1999) Effects of lexicality, frequency, and spelling-to-sound consistency on the functional anatomy of reading. Neuron 24:205-218.

Frankland SM, Greene JD (2015) An architecture for encoding sentence meaning in left mid-superior temporal cortex. Proceedings of the National Academy of Sciences 112:11732-11737.

Friederici AD (2011) The brain basis of language processing: From structure to function. Physiological Reviews 91:1357-1392.

Friederici AD (2012) The cortical language circuit: From auditory perception to sentence comprehension. Trends in Cognitive Sciences 16:262-268.

Goldberg AE (1995) Constructions: A construction grammar approach to argument structure. Chicago, IL: University of Chicago Press.

Gough PM, Nobre AC, Devlin JT (2005) Dissociating linguistic processes in the left inferior frontal cortex with transcranial magnetic stimulation. The Journal of Neuroscience 25:8010-8016.

Grodzinsky Y (2010) The picture of the linguistic brain: How sharp can it be? Reply to Fedorenko & Kanwisher. Language and Linguistics Compass 4:605-622.

Grodzinsky Y, Santi A (2008) The battle for Broca's region. Trends in Cognitive Sciences 12:474-480.

Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R (2004) Intersubject synchronization of cortical activity during natural vision. Science 303:1634-1640.

Hauk O, Davis MH, Pulvermüller F (2008) Modulation of brain activity by multiple lexical and word form variables in visual word recognition: A parametric fMRI study. NeuroImage 42:1185-1195.

Hillis AE (2007) Aphasia progress in the last quarter of a century. Neurology 69:200-213.

Jackendoff R (2002) Foundation of Language: Brain, Meaning, Grammar, Evolution. Oxford, UK: Oxford University Press.

Jackendoff R (2007) A parallel architecture perspective on language processing. Brain Research 1146:2-22.

January D, Trueswell JC, Thompson-Schill SL (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 21:2434-2444.

Joshi AK, Levy LS, Takahashi M (1975) Tree adjunct grammars. Journal of Computer and System Sciences 10:136-163.

Lerner Y, Honey CJ, Silbert LJ, Hasson U (2011) Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. The Journal of Neuroscience 31:2906-2915.

Lescure J (1986) A brief history of OULIPO. In: OULIPO: A Primer of Potential Literature (Motte WF, ed). Lincoln, NE: University of Nebraska Press.

Luria AR, Tsvetkova LS, Futer DS (1965) Aphasia in a composer. Journal of the Neurological Sciences 2:288-292.

McMillan CT, Clark R, Gunawardena D, Ryant N, Grossman M (2012) fMRI evidence for strategic decision-making during resolution of pronoun reference. Neuropsychologia 50:674-687.

McMillan CT, Coleman D, Clark R, Liang T-W, Gross RG, Grossman M (2013) Converging evidence for the processing costs associated with ambiguous quantifier comprehension. Frontiers in psychology 4:153.

Mesulam MM, Thompson CK, Weintraub S, Rogalski EJ (2015) The Wernicke conundrum and the anatomy of language comprehension in primary progressive aphasia. Brain 138:2423-2437.

Nakic M, Smith BW, Busis S, Vythilingam M, Blair RJR (2006) The impact of affect and frequency on lexical decision: The role of the amygdala and inferior frontal cortex. NeuroImage 31:1752-1761.

Nieuwenhuis S, Forstmann BU, Wagenmakers E-J (2011) Erroneous analyses of interactions in neuroscience: A problem of significance. Nature Neuroscience 14:1105-1107.

Nieuwland MS, Petersson KM, Van Berkum JJA (2007) On sense and reference: Examining the functional neuroanatomy of referential processing. NeuroImage 37:993-1004.

Norman-Haignere S, Kanwisher N, McDermott JH (2015) Distinct cortical pathways for music and speech revealed by hypothesis-free voxel decomposition. Neuron 88:1281-1296.

Novais-Santos S, Gee J, Shah M, Troiani V, Work M, Grossman M (2007) Resolving sentence ambiguity with planning and working memory resources: Evidence from fMRI. NeuroImage 37:361-378.

Ojemann G, Ojemann J, Lettich E, Berger M (1989a) Cortical language localization in left, dominant hemisphere: An electrical stimulation mapping investigation in 117 patients. Journal of Neurosurgery 71:316-326.

Ojemann GA, Fried I, Lettich E (1989b) Electrocorticographic (ECoG) correlates of language I: Desynchronization in temporal language cortex during object naming. Electroencephalography and Clinical Neurophysiology 73:453-463.

Peelle JE, Troiani V, Wingfield A, Grossman M (2009) Neural processing during older adults' comprehension of spoken sentences: Age differences in resource allocation and connectivity. Cerebral Cortex 20:bhp142.

Pollard C, Sag IA (1994) Head-driven phrase structure grammar. Chicago, IL: University of Chicago Press.

Rodd JM, Davis MH, Johnsrude IS (2005) The neural mechanisms of speech comprehension: fMRI studies of semantic ambiguity. Cerebral Cortex 15:1261-1269.

Saxe R, Brett M, Kanwisher N (2006) Divide and conquer: A defense of functional localizers. NeuroImage 30:1088-1096.

Schabes Y, Abeille A, Joshi AK (1988) Parsing strategies with "lexicalized" grammars: Application to tree adjoining grammars. In: Proceedings of the 12th Conference on Computational Linguistics, pp 578-583. Stroudsburg, PA.

Stowe LA, Broere CAJ, Paans AMJ, Wijers AA, Mulder G, Vaalburg W, Zwarts F (1998) Localizing components of a complex task: Sentence processing and working memory. Neuroreport 9:2995-2999.

Stromswold K, Caplan D, Alpert N, Rauch S (1996) Localization of syntactic comprehension by positron emission tomography. Brain and language 52:452-473.

Towle VL, Yoon H-A, Castelle M, Edgar JC, Biassou NM, Frim DM, Spire J-P, Kohrman MH (2008) ECoG gamma activity during a language task: Differentiating expressive and receptive speech areas. Brain 131:2013-2027.

Varley RA, Siegal M (2000) Evidence for cognition without grammar from causal reasoning and "theory of mind" in an agrammatic aphasic patient. Current Biology 10:723-726.

Varley RA, Klessinger NJC, Romanowski CAJ, Siegal M (2005) Agrammatic but numerate. Proceedings of the National Academy of Sciences 102:3519-3524.

Wernicke C (1874/1969) The symptom complex of aphasia: A psychological study on an anatomical basis. In: Boston Studies in the Philosophy of Science (Cohen RS, Wartofsky MW, eds), pp 34-97. Dordrecht, The Netherlands: D. Reidel Publishing.

Willems RM, Benn Y, Hagoort P, Toni I, Varley R (2011) Communicating without a functioning language system: Implications for the role of language in mentalizing. Neuropsychologia 49:3130-3135.

Wilson SM, Saygin AP (2004) Grammaticality judgment in aphasia: Deficits are not specific to syntactic structures, aphasic syndromes, or lesion sites. Journal of Cognitive Neuroscience 16:238-252.

Wittgenstein L (1953/2010) Philosophical Investigations. Oxford, UK: Wiley-Blackwell.

Yarkoni T, Speer NK, Balota DA, McAvoy MP, Zacks JM (2008) Pictures of a thousand words: Investigating the neural mechanisms of reading with extremely rapid event-related fMRI. NeuroImage 42:973-987.

Ye Z, Zhou X (2009) Conflict control during sentence comprehension: fMRI evidence. NeuroImage 48:280-290.