Reworking the language network

Evelina Fedorenko[1] and **Sharon L. Thompson-Schill**[2]

[1]Massachusetts Institute of Technology

[2]University of Pennsylvania

Abstract

Prior investigations of functional specialization have focused on the response profiles of particular brain regions. Given the growing emphasis on regional covariation, we propose to reframe these questions in terms of brain "networks" (collections of regions jointly engaged by some mental process). In spite of the challenges that investigations of the language network face, a network approach may prove useful in understanding the cognitive architecture of language. We propose that a language network plausibly includes a functionally specialized "core" (brain regions that coactivate with each other during language processing), and a domain-general "periphery" (a set of brain regions that may coactivate with the language core regions sometimes, but with other specialized systems at other times, depending on task demands). Framing the debate around network properties such as this may prove to be a more fruitful way to advance our understanding of the neurobiology of language.

Keywords

Domain-specificity; domain-generality; language network; cognitive control; fMRI

In search of the language organ

Many deeply important questions in cognition hinge on whether two mental processes rely on the same pool of cognitive and neural resources. Is processing faces distinct from processing other classes of visual objects? Do we use the same mechanisms to extract meaning from words versus pictures? Does resolving linguistic ambiguity draw on the same resources as other demanding tasks? This is one class of questions where fMRI can inform and constrain cognitive theories (cf. [1]; see [2] for discussion). And one such question that has motivated research for decades concerns the existence of a specialized "language organ" (and its possible computations [3]). In particular, are some computations unique to human language, or can language be "solved" by more general-purpose mental operations? Although fMRI cannot directly answer questions of this sort, knowing under which

Corresponding authors:evelina9@mit.edu, sschill@psych.upenn.edu.

conditions a region or a collection of regions is engaged will constrain hypotheses about the likely computations those regions perform.

Traditionally, questions about functional specialization have been asked at the level of individual brain regions (e.g., does Broca's area selectively support speech production?). Techniques, like fMRI, that were lauded for their ability to track regionally-specific changes in metabolic activity seemed perfectly suited to tackle such questions. Today, however, fMRI data are routinely used to describe the statistical interdependencies among brain regions. In particular, hundreds of studies have now reported regional covariation that manifests in a wide range of dependent measures: including signal amplitude—either static, averaged across long time windows of task performance or rest (e.g., *functional connectivity* [4]) or dynamic, varying on a shorter timescale [5]—or, more recently, pattern separability (*informational connectivity* [6]). Because regions that share functional properties can be distant, spanning lobes and hemispheres, their collections are referred to as *large-scale distributed neural networks*, where the regions are the *nodes,* and the inter-regional (implied) connections are the *edges* (Box 1). Furthermore, given that complex cognitive processes – be it face recognition or sentence comprehension – recruit a host of different brain regions [7], it may be time (some might argue, long past time) to start thinking about functional specialization at the level of brain networks (e.g., is the collection of regions recruited by sentence comprehension specialized for solving this particular problem?)

Tests of functional specificity of a brain region routinely begin by defining, on anatomical and/or functional grounds, a region of interest (ROI), and then measuring the relative response of that region across varying cognitive demands (e.g., task conditions). In principle, the logic of assessing the specificity of a brain network is similar: how does a network of interest (NOI) respond across varying cognitive demands? However, in practice, the functional specialization of a network can be evaluated in a number of ways. We here briefly discuss how the notion of functional specialization can be scaled up from brain regions to brain networks, and then we consider the question of functional specialization of the "language network". To foreshadow the take-home message, we argue that the language network includes both relatively specialized and highly domain-general components, and that investigating the dynamic interactions between the two can inform the computations carried out by each.

Scaling up the notion of functional specialization: from nodes to networks

What does it mean for a collection of brain regions to be functionally specialized for some mental process *x*? There are at least three ways to approach this question.

One strategy is to focus on the functional profiles of the individual *nodes*. For example, a network may be functionally specialized for mental process *x* if all of its nodes are functionally specialized for *x* (e.g., Fig. 1a). Or perhaps the presence of at least one functionally specialized node is sufficient for qualifying the whole network as being functionally specialized. (Note that the presence of at least one domain-general node cannot be sufficient for qualifying the network as being domain-general if we are to preserve any

notion of functional specialization, because domain-general processes like attention or cognitive control likely play a role in all mental processes.)

Another strategy is to focus on the *edges* (i.e., the patterns of "connections" among brain regions; cf. Box 1). In this approach, the properties of the nodes are less important: they may be functionally specialized, domain-general, or a mixture of the two. What matters is whether a *unique combination* of nodes and edges is recruited for the relevant mental process *x*. If so, then such a network would be considered functionally specialized for *x*, even if all the individual nodes are domain-general (Fig. 1b); and even the same exact combination of nodes can contribute differently to different mental processes when the nodes are characterized by different patterns of connections.

Yet another possibility is to take time into account and focus on the so-called "network dynamics", i.e. the changes – or lack thereof – in the patterns of coupling (coactivation) between each node and the other nodes of the network as well as the rest of the brain. In particular, networks are not static: they change under different task conditions as well as during different stages of a single task. For example, Cole and colleagues [8] demonstrated that a fronto-parietal brain network shifted its correlation patterns more than did other functional networks as task demands changed. In other words, this fronto-parietal network was promiscuous, partnering with a visual network at times and an auditory network at other times (e.g., Fig. 1d; see also [9,10]). Bassett and colleagues [11] measured how often nodes changed teams across a learning task, and found that the flexibility of a node's team allegiance across time (across individuals) predicted behavior. So the notion of functional specialization may be linked to the stability of the node/network: a network may be functionally specialized for some mental process if its nodes are stable community members, and it may be domain-general if its nodes frequently change allegiances.

The Language "Network of Interest"

We turn now to our primary topic, namely, the language network. In order to ask questions about the nodes, edges, or dynamics of the language network, as described above, we need to define the language network. Immediately, we have a problem: What is language? That is, whatever task (or task comparison) one might choose to define the language network will require assumptions about what putative operations compose language. One could rightly question whether it even makes sense to ask about *a* language network, which presupposes that language is a natural kind (Box 2).

We have two observations about this potential quagmire. Firstly, the daunting task of specifying exactly what language is (and what language isn't) has not completely halted progress: the term "the language network" is being used increasingly frequently (e.g., in the PubMed database an average of 5 articles a year used this term between 2001 and 2005, an average of 17 articles a year between 2006 and 2010, and an average of 35 articles a year since 2011). And secondly, those who have attempted to characterize the language network have not arrived at the same answer to the question of how to define it. Below we consider three common approaches (see Fig. 2, for schematic illustrations of each).

1) Response (e.g., a change in the magnitude of the BOLD signal) to language above a low-level baseline

At first glance, this criterion seems the least controversial: regions that support linguistic processing should respond to language stimuli (e.g., words or sentences). However, although this criterion does include the "classic" regions on the lateral surfaces of the left frontal and temporal cortices, it also includes the regions of the bilateral domain-general cognitive control network (also known as the "task-positive network", the "fronto-parietal attention network", or the "multiple demand network"). The latter network spans extended portions of the frontal and parietal cortices and its regions are engaged in a wide range of goal-directed behaviors [12,13]. Furthermore, for a language *comprehension* task, we would expect to see activity in the primary and higher-level sensory regions: in the auditory cortices for listening [14], and in the visual cortices for reading [15]. And for a language *production* task, we would expect to see activity in the regions that support articulation [16]. In addition, we may observe responses in brain regions implicated in social cognition, emotional processing, high-level visual processing, etc. In the extreme then, the whole brain is probably engaged – in some way – during language processing, but the notion of a language network is only useful if we agree that some subset of the brain is more strongly, consistently, and/or causally engaged in language processing than the rest of the brain.

2) Response to a conjunction of either listening and reading, or comprehension and production, above a low-level baseline

These criteria will lead to the elimination of the sensory regions [17,18] and/or the articulation regions [19]. However, at least some parts of the cognitive control network are still likely to be included.

3) A greater response to intact language stimuli than to various degraded versions of those stimuli, matched for low-level features

This criterion is frequently used and translates into contrasts like speech vs. backwards speech [20], sentences vs. false fonts [21], sentences vs. lists of unconnected words [17,22], including parametric versions of this contrast [23], sentences vs. lists of nonwords [17], etc. Depending on how low-level the baseline condition is, these contrasts may or may not include parts of the higher-level sensory regions (e.g., the visual word-form area for visual stimuli). Importantly though, these contrasts will generally *not* include the cognitive control regions because those regions tend to respond more strongly to linguistically degraded conditions [13,24], presumably because of the greater cognitive effort required to process degraded stimuli (cf. similar effects in visual processing where non-degraded images produce stronger responses in the visual cortices, but degraded images more strongly activate the cognitive control network [25]).

Of course, a particular functional signature (as in 1-3 above; Fig. 2) is not the only way to define the language network. For example, one might want to emphasize the *causal* role in linguistic processing, such that a region is included only if a transient or more permanent disruption of its activity leads to linguistic deficits. Another possible criterion has been adopted by Hasson and colleagues who have focused on the across-subjects similarity in the timecourses of neural activity. For example, in studies of narrative comprehension, they

reported inter-subject correlations across large extents of cortex, including what-appear-to-be cognitive control regions in the dorsolateral prefrontal cortex and around inferior parietal sulcus [26]. And they further showed that even some regions that do not show an above-baseline response to language (cf. criterion 1), like the precuneus, show correlated activity across subjects during sentence comprehension [27].

So where does this leave us? It should be clear that the "right" criterion (or set of criteria) for the language network is subject to debate. And, the multitude of possible definitions of the language network – especially differences in whether the domain-general cognitive control regions are included – has almost certainly contributed to past disagreements, some of which may have been superficial (definitional) in nature. But this problem is neither unique to the domain of language, nor to the network-level approach, similarly characterizing investigations that have focused on single brain regions. And although any conclusions one draws about a brain network are bound to be affected by how that network is defined, we cannot abandon the enterprise simply because we cannot agree on the inclusion criteria. Instead, as we will argue below, network approaches may, in fact, be able to help us sharpen the definition of the language network, which will, in turn, have implications for understanding the specificity of the cognitive and computational mechanisms required during language production and comprehension.

Is the language network functionally specialized?

The extent to which language – including its many components (like speech perception, letter/word recognition, articulation, syntactic processing, etc.; Box 2) – relies on functionally specialized vs. domain-general cognitive and neural machinery has been long debated. One important take-away message from the preceding section is that under many definitions, the language network includes both relatively functionally specialized brain regions [24,28–30] and brain regions better thought of as part of a domain-general cognitive control network [31–35] (for a review see [36]). But, our opening question, and a central concern for the field, is whether *language* requires specific computations, not to characterize any particular brain region. That is, is the language network – which comprises both specialized and domain-general nodes – functionally specialized for language processing (or some aspects thereof)? Let us revisit the possible definitions of a functionally specialized brain network sketched out above.

If we focus on the properties of the individual *nodes*, then one could argue either that a) the language network is not functionally specialized for language because not all of its nodes are functionally specialized for language; or that b) the language network is functionally specialized, if the presence of *some* specialized nodes is sufficient.

If we focus on the *edges*, the language network would qualify as functionally specialized because the specialized regions would only get engaged during language processing tasks and thus, by definition, the *combination* of brain regions (and presumably the connections among them) engaged during language processing would be unique, as no other mental process would recruit the specialized regions of the language network.

Finally, let us consider the dynamic aspects of the language network. It is important to note that although the functional profiles of the different brain regions comprising this network vary considerably – with some being relatively functionally specialized for language processing and others highly domain-general – these sets of brain regions must nevertheless have a way to communicate with one another (cf. Fig. 1a). Such interactions are essential given that domain-general brain regions – like those that support cognitive control, working memory, attention, etc. – likely participate in all mental processes, including language comprehension and production. And the role of these domain-general mechanisms in language processing is hard to dispute given the abundant behavioral evidence from dual-task paradigms [37–39] and individual-differences investigations [40–45], as well as neuroimaging evidence from many linguistic complexity manipulations [46–50]. Thus any putatively language-specific components of the language network cannot be encapsulated in the Fodorian sense (cf. Fig. 1a).

But how exactly the functionally specialized and the domain-general components of the language network interact remains to be discovered. For example, such interactions could take place via stable hubs [8] that are always "partnering" with one or more specialized networks (Fig. 1c). Alternatively, they could occur via dynamic changes in the patterns of connections between the specialized and domain-general sets of regions (Fig. 1d). For example, during language processing, the language network configuration may look like Fig. 1d at t=1, with the specialized (pink) and domain-general (multi-colored) regions working together. At other times, the same domain-general regions may partner with a different network (e.g., Fig. 1d at t=2).

This notion of node "stability" has motivated a distinction between a network *core* (i.e., a set of brain regions that – consistently, across time – interact with one another) and a network *periphery* (i.e., a set of brain regions that interact with a different specialized set of brain regions at different times, depending on current goals [51]). We propose that this distinction between the core and periphery of a network could be a useful tool for exploring the functional specialization of any network, including the language network. In our review of this question above and in our previous work [28,36], we focused on cross-task structure (i.e., functionally specific regions are those that are active in specific tasks but not across tasks whereas domain-general regions are active across tasks). In dynamic network terms, the focus is on cross-time structure (core regions are those that maintain their allegiance through time whereas peripheral regions do not). These two dimensions are plausibly not independent. In particular, the "promiscuity" of a brain region (i.e., its tendency to couple – i.e., coactivate – with other brain regions) may be inversely related to the degree of its functional specialization. In particular, if a brain region supports computations that are well-suited for solving a particular task (including, for example, relying on domain-specific knowledge structures), it will likely only partner with regions that implement similar, specific, computations. In contrast, if a brain region supports highly generic computations (e.g., exclusive OR [52]), it can partner with a wide array of functionally diverse regions because such generic computations are plausibly useful across domains. Thus, a language network plausibly includes a functionally specialized core, and a domain-general periphery. And the existence of core regions supporting a particular mental process is perfectly

compatible with the importance of domain-general circuits for that mental process (cf. [53]; see [54,55] for discussions of the relationship between functional specificity and encapsulation).

As we foreshadowed above, this dynamic network approach may enable one to *discover* the language network a) without needing to define it *a priori,* or b) starting with a very liberal, all inclusive, definition. Any given language task (e.g., sentence comprehension) is likely to require both domain-general processes and (if they exist) language-specific processes; however, assuming those processes are supported by different nodes in a network, the dynamic relations among these nodes should vary across the task in a way that is detectable with network approaches and that potentially informs the function of the nodes. This, in turn, can motivate future studies aimed at understanding the response properties of different sets of nodes (whether they be "core" or "periphery" in one task) under different cognitive conditions. And so we proceed.

Concluding remarks

For many years, we—both we, the individual scientists co-authoring this piece and we, two representative researchers of the neurobiology of language—have been trying to understand the cognitive and neural architecture of language using regionally-specific fMRI responses. Much of this work has been framed as a debate about the functional specificity of regions recruited during language processing, and we have contributed to each side of this debate. One of us has argued that there is "a high degree of functional specificity in the brain regions that support language" [28]. The other has advanced hypotheses about putative language regions that are, instead, "grounded in domain-general terms*"* [41]. Here, we have joined forces in an attempt to redirect the empirical enterprise, by calling a ceasefire on arguments about whether an individual brain region is domain-specific or domain-general. Instead, we propose that our understanding of the computations that enable language and the neural systems that support them will advance more rapidly if we follow the example of many other subfields of neuroscience and turn our efforts towards characterizing properties of the language network (or networks). We have outlined here a strategy for characterizing some properties of these networks, including properties that may speak directly to the specificity of language functions. We have argued that this approach does not require agreement on how to define language in order to begin to characterize properties of these networks. We have borrowed the concepts of a network "core" and "periphery" from other areas of network neuroscience, in order to propose a means for identifying functionally-specific and functionally-diverse nodes in a language network, by tracking the structure of their responses across time (instead of across tasks). We are confident that this approach will identify both domain-specific and domain-general machinery, and we can then begin to understand the relative contributions of each, as well as the dynamics of their interaction and how those dynamics affect behavioral outcomes (Box 3).

Acknowledgments

References

1. Coltheart M. How can functional neuroimaging inform cognitive theories? Perspect. Psychol. Sci. 2013; 8:98–103.

2. Mather M, et al. How fMRI can inform cognitive theories. Perspect. Psychol. Sci. 2013; 8:108–113. [PubMed: 23544033]

3. Hauser MD, et al. The faculty of language: what is it, who has it, and how did it evolve? Science. 2002; 298:1569–1579. [PubMed: 12446899]

4. Friston KJ. Functional and effective connectivity: a review. Brain Connect. 2011; 1:13–36. [PubMed: 22432952]

5. Hutchison RM, et al. Dynamic functional connectivity: promises, issues, and interpretations. NeuroImage. 2013; 80:360–378. [PubMed: 23707587]

6. Coutanche MN, Thompson-Schill SL. Informational connectivity: identifying synchronized discriminability of multi-voxel patterns across the brain. Front. Hum. Neurosci. 2013; 7:15. [PubMed: 23403700]

7. Mesulam M. Large-scale neurocognitive networks and distributed processing for attention, language, and memory. Ann. Neurol. 1990; 28:597–613. [PubMed: 2260847]

8. Cole MW, et al. Multi-task connectivity reveals flexible hubs for adaptive task control. Nat. Neurosci. 2013; 16:1348–1355. [PubMed: 23892552]

9. Chadick JZ, Gazzaley A. Differential coupling of visual cortex with default or frontal-parietal network based on goals. Nat. Neurosci. 2011; 14:830–832. [PubMed: 21623362]

10. Al-Aidroos N, et al. Top-down attention switches coupling between low-level and high-level areas of human visual cortex. Proc. Natl. Acad. Sci. U. S. A. 2012; 109:14675–14680. [PubMed: 22908274]

11. Bassett DS, et al. Dynamic reconfiguration of human brain networks during learning. Proc. Natl. Acad. Sci. U. S. A. 2011; 108:7641–7646. [PubMed: 21502525]

12. Duncan J. The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. Trends Cogn. Sci. 2010; 14:172–179. [PubMed: 20171926]

13. Fedorenko E, et al. Broad domain generality in focal regions of frontal and parietal cortex. Proc. Natl. Acad. Sci. U. S. A. 2013; 110:16616–16621. [PubMed: 24062451]

14. DeWitt I, Rauschecker JP. Wernicke's area revisited: parallel streams and word processing. Brain Lang. 2013; 127:181–191. [PubMed: 24404576]

15. Yeatman JD, et al. Development of white matter and reading skills. Proc. Natl. Acad. Sci. 2012; 109:E3045–E3053. [PubMed: 23045658]

16. Grabski K, et al. Functional MRI assessment of orofacial articulators: neural correlates of lip, jaw, larynx, and tongue movements. Hum. Brain Mapp. 2012; 33:2306–2321. [PubMed: 21826760]

17. Fedorenko E, et al. New method for fMRI investigations of language: defining ROIs functionally in individual subjects. J. Neurophysiol. 2010; 104:1177–1194. [PubMed: 20410363]

18. Braze D, et al. Unification of sentence processing via ear and eye: an fMRI study. Cortex J. Devoted Study Nerv. Syst. Behav. 2011; 47:416–431.

19. Menenti L, et al. Shared language: overlap and segregation of the neuronal infrastructure for speaking and listening revealed by functional MRI. Psychol. Sci. 2011; 22:1173–1182. [PubMed: 21841148]

20. Bedny M, et al. Language processing in the occipital cortex of congenitally blind adults. Proc. Natl. Acad. Sci. 2011; 108:4429–4434. [PubMed: 21368161]

21. Noppeney U, Price CJ. Retrieval of abstract semantics. NeuroImage. 2004; 22:164–170. [PubMed: 15110006]

22. Snijders TM, et al. Retrieval and unification of syntactic structure in sentence comprehension: an FMRI study using word-category ambiguity. Cereb. Cortex N. Y. N 1991. 2009; 19:1493–1503.

23. Pallier C, et al. Cortical representation of the constituent structure of sentences. Proc. Natl. Acad. Sci. U. S. A. 2011; 108:2522–2527. [PubMed: 21224415]

24. Fedorenko E, et al. Language-selective and domain-general regions lie side by side within Broca's area. Curr. Biol. CB. 2012; 22:2059–2062.

25. Gorlin S, et al. Imaging prior information in the brain. Proc. Natl. Acad. Sci. 2012; 109:7935–7940. [PubMed: 22538820]

26. Lerner Y, et al. Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. J. Neurosci. Off. J. Soc. Neurosci. 2011; 31:2906–2915.

27. Ben-Yakov A, et al. Loss of reliable temporal structure in event-related averaging of naturalistic stimuli. NeuroImage. 2012; 63:501–506. [PubMed: 22813575]

28. Fedorenko E, et al. Functional specificity for high-level linguistic processing in the human brain. Proc. Natl. Acad. Sci. U. S. A. 2011; 108:16428–16433. [PubMed: 21885736]

29. Monti MM, et al. The boundaries of language and thought in deductive inference. Proc. Natl. Acad. Sci. 2009; 106:12554–12559. [PubMed: 19617569]

30. Monti MM, et al. Thought beyond language: neural dissociation of algebra and natural language. Psychol. Sci. 2012; 23:914–922. [PubMed: 22760883]

31. Thompson-Schill SL, et al. Role of left inferior prefrontal cortex in retrieval of semantic knowledge: a reevaluation. Proc. Natl. Acad. Sci. U. S. A. 1997; 94:14792–14797. [PubMed: 9405692]

32. Thompson-Schill SL, et al. Effects of repetition and competition on activity in left prefrontal cortex during word generation. Neuron. 1999; 23:513–522. [PubMed: 10433263]

33. Bedny M, et al. Semantic adaptation and competition during word comprehension. Cereb. Cortex N. Y. N 1991. 2008; 18:2574–2585.

34. Hindy NC, et al. The effect of object state-changes on event processing: do objects compete with themselves? J. Neurosci. Off. J. Soc. Neurosci. 2012; 32:5795–5803.

35. Hindy NC, et al. A cortical network for the encoding of object change. Cereb. Cortex N. Y. N 1991. 2013 DOI: 10.1093/cercor/bht275.

36. Thompson-Schill SL, et al. The frontal lobes and the regulation of mental activity. Curr. Opin. Neurobiol. 2005; 15:219–224. [PubMed: 15831406]

37. Gordon PC, et al. Memory-load interference in syntactic processing. Psychol. Sci. 2002; 13:425–430. [PubMed: 12219808]

38. Fedorenko E, et al. The nature of working memory in linguistic, arithmetic and spatial integration processes. J. Mem. Lang. 2007; 56:246–269.

39. Rodd JM, et al. The role of domain-general frontal systems in language comprehension: evidence from dual-task interference and semantic ambiguity. Brain Lang. 2010; 115:182–188. [PubMed: 20709385]

40. Carretti B, et al. Role of working memory in explaining the performance of individuals with specific reading comprehension difficulties: A meta-analysis. Learn. Individ. Differ. 2009; 19:246–251.

41. Novick JM, et al. A case for conflict across multiple domains: memory and language impairments following damage to ventrolateral prefrontal cortex. Cogn. Neuropsychol. 2009; 26:527–567. [PubMed: 20183014]

42. Cragg L, Nation K. Language and the development of cognitive control. Top. Cogn. Sci. 2010; 2:631–642.

43. Khanna MM, Boland JE. Children's use of language context in lexical ambiguity resolution. Q. J. Exp. Psychol. 2006. 2010; 63:160–193.

44. McVay JC, Kane MJ. Why does working memory capacity predict variation in reading comprehension? On the influence of mind wandering and executive attention. J. Exp. Psychol. Gen. 2012; 141:302–320. [PubMed: 21875246]

45. Novick JM, et al. Clearing the garden-path: improving sentence processing through cognitive control training. Lang. Cogn. Process. 2013:1–44. 0.

46. Novais-Santos S, et al. Resolving sentence ambiguity with planning and working memory resources: evidence from fMRI. NeuroImage. 2007; 37:361–378. [PubMed: 17574445]

47. January D, et al. Co-localization of stroop and syntactic ambiguity resolution in Broca's area: implications for the neural basis of sentence processing. J. Cogn. Neurosci. 2009; 21:2434–2444. [PubMed: 19199402]

48. McMillan CT, et al. fMRI evidence for strategic decision-making during resolution of pronoun reference. Neuropsychologia. 2012; 50:674–687. [PubMed: 22245014]

49. Nieuwland MS, et al. Brain regions that process case: evidence from Basque. Hum. Brain Mapp. 2012; 33:2509–2520. [PubMed: 21898678]

50. Wild CJ, et al. Effortful listening: the processing of degraded speech depends critically on attention. J. Neurosci. Off. J. Soc. Neurosci. 2012; 32:14010–14021.

51. Bassett DS, et al. Task-based core-periphery organization of human brain dynamics. PLoS Comput. Biol. 2013; 9:e1003171. [PubMed: 24086116]

52. Rigotti M, et al. Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. Front. Comput. Neurosci. 2010; 4:24. [PubMed: 21048899]

53. Blumstein SE, Amso D. Dynamic functional organization of language: insights from functional neuroimaging. Perspect. Psychol. Sci. 2013; 8:44–48.

54. Coltheart M. Modularity and cognition. Trends Cogn. Sci. 1999; 3:115–120. [PubMed: 10322463]

55. Barrett HC, Kurzban R. Modularity in cognition: framing the debate. Psychol. Rev. 2006; 113:628–647. [PubMed: 16802884]

56. Honey CJ, et al. Predicting human resting-state functional connectivity from structural connectivity. Proc. Natl. Acad. Sci. 2009; 106:2035–2040. [PubMed: 19188601]

57. O'Reilly JX, et al. Causal effect of disconnection lesions on interhemispheric functional connectivity in rhesus monkeys. Proc. Natl. Acad. Sci. 2013; 110:13982–13987. [PubMed: 23924609]

58. Le Bihan D, et al. Artifacts and pitfalls in diffusion MRI. J. Magn. Reson. Imaging. 2006; 24:478–488. [PubMed: 16897692]

59. Jones DK, Cercignani M. Twenty-five pitfalls in the analysis of diffusion MRI data. NMR Biomed. 2010; 23:803–820. [PubMed: 20886566]

60. Jones DK, et al. White matter integrity, fiber count, and other fallacies: the do's and don'ts of diffusion MRI. Neuroimage. 2013

61. Yeo BTT, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J. Neurophysiol. 2011; 106:1125–1165. [PubMed: 21653723]

62. Power JD, et al. Functional network organization of the human brain. Neuron. 2011; 72:665–678. [PubMed: 22099467]

63. Bassett DS, Lynall ME. Network methods to characterize brain structure and function. 2013 Manuscript submitted for publication.

64. Baker CI, et al. Visual word processing and experiential origins of functional selectivity in human extrastriate cortex. Proc. Natl. Acad. Sci. U. S. A. 2007; 104:9087–9092. [PubMed: 17502592]

65. Hamamé CM, et al. Dejerine's reading area revisited with intracranial EEG: selective responses to letter strings. Neurology. 2013; 80:602–603. [PubMed: 23382370]

66. Bohland JW, Guenther FH. An fMRI investigation of syllable sequence production. NeuroImage. 2006; 32:821–841. [PubMed: 16730195]

67. Goldberg AE. Constructions: a new theoretical approach to language. Trends Cogn. Sci. 2003; 7:219–224. [PubMed: 12757824]

68. Jackendoff, R. Foundations of language: Brain, meaning, grammar, evolution. Oxford University Press; 2002.

69. Bybee, JL. Language, usage and cognition. Cambridge University Press; Cambridge: 2010.

70. Gibson E. Linguistic complexity: locality of syntactic dependencies. Cognition. 1998; 68:1–76. [PubMed: 9775516]

71. McElree B, et al. Memory structures that subserve sentence comprehension. J. Mem. Lang. 2003; 48:67–91.

72. Gordon PC, et al. Memory interference during language processing. J. Exp. Psychol. Learn. Mem. Cogn. 2001; 27:1411–1423. [PubMed: 11713876]

73. Lewis RL, et al. Computational principles of working memory in sentence comprehension. Trends Cogn. Sci. 2006; 10:447–454. [PubMed: 16949330]

74. Hale, J. A probabilistic Earley parser as a psycholinguistic model. Proceedings of the second meeting of the North American Chapter of the Association for Computation Linguistics on Language Technologies; Stroudsburg, PA, USA. 2001. p. 1-8.

75. Levy R. Expectation-based syntactic comprehension. Cognition. 2008; 106:1126–1177. [PubMed: 17662975]

76. Dehaene S, Changeux JP. Experimental and theoretical approaches to conscious processing. Neuron. 2011; 70:200–227. [PubMed: 21521609]

BOX 1

Challenges for network neuroscience

a. The requisite care in using the term "network"

Although the terms "network" and "connectivity" are widely used when talking about regional covariation in the human brain, it is important to keep in mind that no human data at present allow us to make inferences about brain regions forming networks in the true sense of the word. In particular, under a technical definition, two brain regions form a network if they are anatomically connected, typically via monosynaptic projections. In living humans, we rarely, if ever, can say anything conclusive about anatomical connections among brain regions. In particular, functional correlation data (task-based or resting state) cannot be used to infer anatomical connectivity because the relationship between the two is complex [56,57], and diffusion tractography is still severely limited [58–60]. Consequently, although we follow the literature in adopting the term "network" to refer to collections of regions that share functional properties, these collections of regions are more appropriately characterized as *functional systems*.

b. Uncertainty about the number and nature of brain networks

Network neuroscience is still in its youth, and there is at present no agreement on many important questions. For example, what is the right way to parcellate the brain into regions (nodes)? How many networks does the human brain encompass, and how to functionally interpret these networks (see [61,62] for some recent proposals)? What is the structure of each network, and does this structure differ across networks (see [63] for a review)? The answer to these questions is likely to arise from a combination of data-driven approaches that attempt to carve up the entire brain based on patterns of inter-regional co-variation and approaches that more narrowly focus on a particular subset of the brain (e.g., the language network), because a deep understanding of the structure and function of one network – including the contributions of each node and edge and their dynamics – may shed light on the broader functional architecture of the human brain.

BOX 2

Is "language" a natural kind?

One might object that questions about the language network are ill-posed, because language is not a single thing. Indeed, when talking about whether "language" relies on domain-specific vs. domain-general machinery (or some combination of the two), researchers are often referring to different mental processes that language encompasses, and there is no agreement on the right ontology of these processes. Such ontologies in human cognitive neuroscience are typically inspired by theoretical and experimental behavioral work in psychology and cognitive science, although often lag behind. At present, based on differences in functional profiles and some neuropsychological patient evidence, we can at least distinguish between i) the sensory language regions (in the auditory and visual cortices), ii) the speech articulation regions, and iii) the "higher-level" language processing regions (Fig. 2). For example, in contrast to the high-level language regions, the sensory regions appear to respond to stimuli that are devoid of meaning: the visual word-form area responds as much to consonant strings as to real words [64,65]. Similarly, the speech articulation regions [66] can be driven by low-level production tasks like repeating nonwords. Even within the "higher-level" language interpretation we may want to distinguish between phonological (sound-level) processing, lexical (word-level) processing, and combinatorial (syntactic and compositional semantic) processing, although more recent, usage-based, approaches argue against sharp boundaries between them [67–69] (see [24] for some fMRI support). In summary, although language processing plausibly relies on several distinct kinds of computations implemented in distinct brain regions, there appears to be a subset of our brain that is consistently engaged when we produce and/or understand language that is at least partially distinct from brain regions that support other mental processes [28,30]. Talking about "language" as a natural kind is thus not unreasonable, although of course if/when dissociations among different language components are discovered – be it by studying the functional response profiles of the relevant regions or their dynamic network properties – those dissociations should be taken into account when thinking about the computations that those components require.

BOX 3

Outstanding questions

- Are any of the regions engaged in language processing truly selective for language?

  Fedorenko and colleagues [24,28] have shown that a number of language regions show functionally specific responses. However, other non-linguistic processes remain to be tested, and either outcome—functional specificity vs. overlap with some non-linguistic processes in some or all of the language regions—will further constrain the space of possibilities for what these regions could be doing and may reveal important differences among regions.

- What role do domain-general cognitive control brain regions play in language comprehension and production?

  As discussed in the main text, the engagement of domain-general cognitive control / working memory mechanisms in language is not under debate, but many important questions remain about *the nature and significance of this engagement*. In particular:

  a. Given the extent (and the bilateral nature) of the cognitive control network, are some of its components more important for language processing than others? And if so, which ones?

  b. What is the significance, if any, of the proximity of the language and cognitive control regions in the left frontal lobe [24]? For example, perhaps some properties of the cells in the left frontal cortex or its position within the large-scale network structure make it well-suited for performing both some aspects of language processing and the kinds of generic computations that are likely to take place in the cognitive control regions [36].

  c. What are the precise circumstances under which cognitive control regions get engaged during language processing? Can we predict their engagement from current models of linguistic complexity (e.g., memory-based models [70–73]; or experience-/surprisal-based models [74,75])?

  d. Do cognitive control regions provide alternative routes for "solving" language, or do they merely provide extra computational power (a "workspace" of sorts [76]) when the core language regions run into difficulty?

  e. What is the time-course of the interaction between the language regions and the cognitive control regions? For example, when an ambiguous word or structure is encountered, is it the language or cognitive control regions that respond first? And how is the interaction between these regions manifested: that is, do they show increases in synchronization

of neural activity (which would suggest parallel – although possibly distinct – computations)? anti-correlations (which would suggest some kind of a trade-off)?

- How might network dynamics vary across individuals, and what might that variation predict about individual differences in behavior?

Highlights

- A central question in cognitive science concerns whether human language requires specialized computational machinery or whether it can be solved by more general-purpose mental operations; in recent years, fMRI data have been used to argue for both sides of this debate.

- The systems that support language processing might be better described at the level of interactive networks, not individual brain regions.

- Many different criteria have been used to identify a "language network," and variations in these definitions may contribute to some of the apparent disagreements in the field.

- Measuring the dynamic properties of neural networks can reveal "core" and "peripheral" components of a network; the cross-time structure of the language network provides a new tool for thinking about functional specificity.
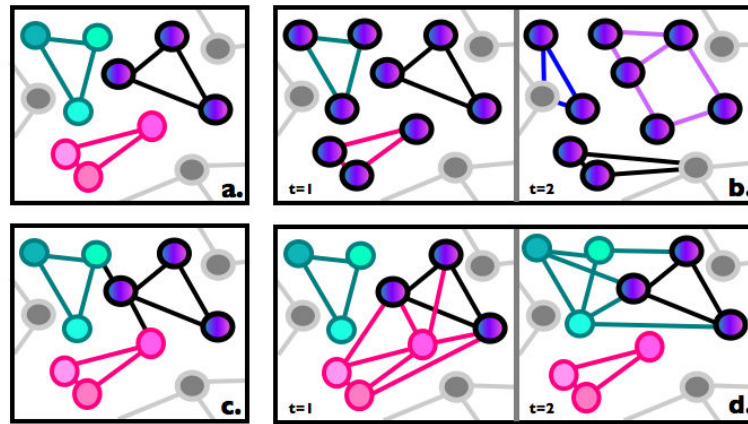
Figure 1. Hypothetical network configurations

A schematic illustration of several hypothetical network configurations (for b and d, we show network configurations at two different time points). Single-color nodes represent functionally specific brain regions (i.e., regions that are selectively engaged by a particular mental process); multiple-color nodes represent domain-general brain regions (i.e., regions with functionally diverse responses). 1a. Two networks composed of functionally specific brain regions, and one network composed of domain-general brain regions, with no interaction between the networks. 1b. No functionally specific brain regions; all brain regions are domain-general but combine in different ways to solve different tasks. 1c. Two networks composed of functionally specific brain regions, and one network composed of domain-general brain regions, with one node of the latter serving as a "hub" via which the specialized networks can interact with the domain-general network. 1d. Two networks composed of functionally specific brain regions, and one network composed of domain-general brain regions. The latter can combine with either of the two specialized networks at different times, thus becoming the "periphery" of the pink network at t=1, and of the green network at t=2.
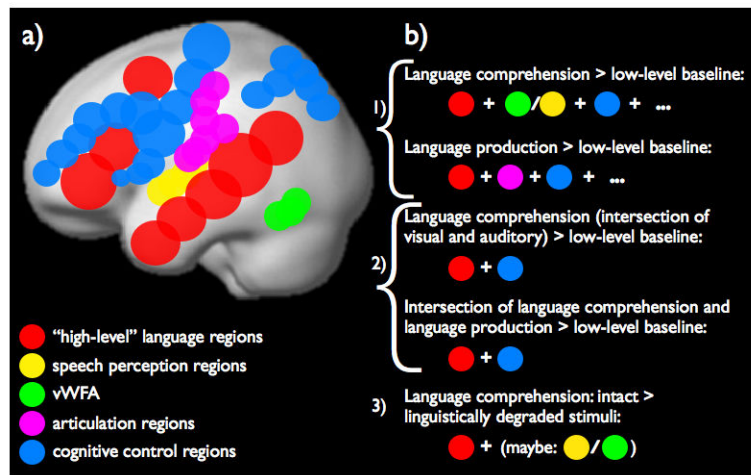
Figure 2. The language network under different definitions
a) A schematic depiction of five sets of brain regions that are sometimes included in the language network: red = the classic high-level language processing regions; yellow = speech perception regions; green = visual word-form area; pink = speech articulation regions; and blue = cognitive control regions. b) A schematic illustration of possible definitions of the language network, ranging from very liberal (1) to more conservative (2 and 3).