

7 A Mosaic of Neurolinguistic Modules

As you grow older, the brain cells connected with anxiety begin to die. . . . You start feeling better and your voice gets deeper.

—Leonard Cohen, BBC interview, 1988

We must not ignore anatomy when speaking of the physical basis of words.

—John Hughlings Jackson, FRS, neurologist, 1878

Gabby Giffords, a Story of Language Loss and of Courage

On January 8, 2011, a lone shooter attacked a small crowd in the parking lot of a shopping center in Tucson, Arizona. His target was Gabby Giffords, a first term Arizona congresswoman, who was holding a meeting with voters. A total of 18 people were shot, six of them fatally, including a nine-year-old girl and a judge. Giffords survived the attack. A bullet went into the rear part of her brain, exiting from its front left side. On its way, the bullet damaged motor and language regions, leaving her partially paralyzed and language-impaired.

Giffords's personal story is one of courage, optimism, and resilience: together with her husband, former astronaut and current US senator Mark Kelly, she has become a fierce and outspoken supporter of gun control legislation. Judging by her media appearances, she remains very active publicly more than a dozen years later, despite her limitations, which seem to be mostly in the motor and language domains.

I have not met or examined Gabby Giffords, nor have I had access to her records. But looking at her powerful public appearances (and especially, having watched the fascinating and uplifting documentary *Gabby Giffords*



REP. GIFFORDS SHOT, CRITICAL

Judge Roll, girl, aide, retiree, 2 others slain

Figure 7.1

A local Arizona headline reporting the assassination attempt.

Won't Back Down),¹ it is reasonably clear to me that she suffers from a language deficit clinically known as Broca's aphasia. A relatively mild patient, her speech is somewhat nonfluent and hesitant, and she makes grammatical errors that are typical of this syndrome. Her overall cognitive state appears intact, yet sentences she produces are short and simple, rarely containing embeddings, and she has difficulty repeating sentences uttered by others. Here is an example, recorded as Giffords was rehearsing her congressional resignation speech in front of a camera: "Action! [I] gotta step down . . . people [of] Arizona . . . step down people [of] Arizona. Right now [there's] work to do . . . work to do . . . work to do, work . . . Oh." Watching her speak a decade later, she seems to have improved a lot, but she is still language-deficient, and she appears to suffer from a motor deficit to her right side as well.

To get an insight into the brain bases of language, I will take Giffords's personal story of resilience and courage as a starting point for the clinical and scientific field of neurolinguistics. In this chapter, I will tell you a bit about this field, which tries to identify and characterize language mechanisms in the human brain. It does so on the basis of clinical observations of brain-damaged patients, functional imaging in the healthy speaking brain, and intraoperative test scores of patients who undergo awake brain

surgery—three research areas in which I have long been deeply involved. I believe that we can understand the present better once we know how we got here, and so I will try to present the current state of the art together with its history.

Paul Broca and His Legacy

Gabby Giffords's cluster of signs is well-recognized, known in the medical literature since French surgeon Paul Broca presented the case of the deceased Mr. Leborgne at a meeting of the Parisian *Société Anatomique* in 1861.² Broca had examined Leborgne, who could hardly utter a word, and taken notes on his general behavior. Subsequent to the patient's death, Broca had inspected his brain. In his presentation, Broca emphasized the selective nature of the patient's deficit, which he described as a severe impairment to his spoken language and not much else. The behavioral limitation was causally related to the "softening of the second and third convolutions of the upper part of the left frontal lobe." This led Broca to claim that he had discovered "the seat of articulated language" (*le siège de la faculté du langage articulé*) in the human brain. Importantly, Broca never dissected Leborgne's brain. Neuroanatomy at his time was not well-developed (most scientists did not even believe that the brain was composed of cells), and his characterization of the anatomy was rather opaque. Broca thus left the brain intact (it rests preserved in a Parisian museum to this day), inspected the lesion to its surface, and stopped short of assessing its precise depth or extent.

His limited empirical observations did not stop Broca from making three bold claims regarding brain structure and function, in opposition to the beliefs of most of his contemporaries: first, *distinctness*—the idea that language constitutes a separate cognitive function; second, *uniqueness*—that this function has a unique brain location, that is, a dedicated piece of neural tissue that supports it; and third, *laterality*—that language is an exclusively left-hemispheric function, from which it follows that the visually apparent anatomic symmetry between the two halves of the brain is not an indicator of functional symmetry. As Broca emphasized, patient Leborgne's loss of language was only due to damage on the left side.

Broca's contribution to science and medicine soon became widely recognized. The brain locus of this patient's lesion was later named Broca's area, and a Paris street was even named after him in 1890; see figure 7.2. Despite



Figure 7.2

A street sign for Rue Broca in Paris. Photograph courtesy of Chabe1 (Creative Commons license CC-BY-SA 4.0).

objections from many quarters, his approach to brain structure and function has guided modern neurology for over 160 years.

Pieces of Grammar in Pieces of Brain

Two left hemispheric regions, Broca's, and one now known as Wernicke's area, have long been considered the main language "areas" of the human brain. As can be seen in figure 7.3 (taken from a seminal 1970 paper, which helped Harvard neurologist Norman Geschwind revive interest in brain-language relations),³ Broca's area (**B**) and Wernicke's area (**W**) are connected by a thick bundle of fibers.

On this classical view, inspired by an 1885 predecessor,⁴ **B** is an area entrusted with the production of language and speech, **W** is in charge of language reception, **A** is a store of words (the mental lexicon), and the cable-like structure connecting **B** and **W**—a bundle of fibers known as the arcuate fasciculus—presumably relays information between these areas.




Figure 7.3

Norman Geschwind's schema of the brain's language map. Reproduced by permission from N. Geschwind, "The Organization of Language and the Brain," *Science* 170, no. 3961 (1970): 940–944.

This schema is, for all I know, the first anatomical map of a higher cognitive function ever (barring earlier phrenological drawings that were made on close to zero empirical evidence):⁵ anatomical locations (large neuronal chunks) are assigned a functional role, while a thick bundle of (white matter) fibers functions as a connector cable. Indeed, this approach became known as connectionism. It was conceived on the basis of observations of patients who had suffered a focal, typically stroke-induced lesion to left Broca's or Wernicke's area, as well as to the arcuate fasciculus. Note that the patients who are most relevant to the present discussion are those with a *unifocal* brain disease; that is, the functional impairment is related to a single, identifiable, and hopefully well-delineated lesion. Patients with multiple strokes or tumors or with neurodegenerative ailments such as Alzheimer's disease do not fit this schema easily, because their pathology is either *multifocal* or *diffuse*, which stands in the way of efforts to achieve precise functional neuroanatomy.

The degree of detail of the old connectionist theory as depicted in figure 7.3 is obviously far from satisfactory. Anatomically, the extent of the named areas is at best vague, as they were marked out in an era that had no sophisticated neuroimaging devices; functionally, the terms *language production* and *language reception* refer to modalities through which language is

practiced, but they make no contact with any current theoretical account of linguistic knowledge or ability. Still, Broca's, Wernicke's, and Geschwind's perspective carries an important message for cognitive neuroscientists: *pieces of cognition reside in pieces of brain*. This is neurolinguistics—the study of the functional anatomy of the network entrusted with linguistic ability in its entirety. Our task is to identify the details of the functional architecture of language and its precise brain bases. Focal brain disease is a good place to start, as it may provide important functional and anatomical evidence. Regarding the tools that may help us to discover this architecture, there are two obvious lead candidates: (psycho)linguistics on the functional side and micro- and macro-anatomy on the structural side.

Functional Modules: Our Auditory Brain Is a Phonetician

The first to understand the importance of the connection between linguistics and the brain was structural linguist Roman Jakobson, who in the 1940s started investigating patterns of phonetic breakdown in the speech of individuals with aphasia.⁶ Earlier, when residing in Prague, he and his fellow Russian linguist Nikolai Trubetzkoy had proposed a theory that phonetic entities, or phonemes, are contrast-based. As we noted in chapter 1, consonants *p*, *t*, *k* are contrasted with *b*, *d*, *g* in that only the latter are “voiced” (i.e., the vocal cords start vibrating 100 milliseconds earlier than while articulating the “unvoiced” counterpart); *p*, *t*, *k* are also contrasted with *ph*, *th*, *kh* in the “manner” of their articulation (plosive vs. fricative, i.e., whether the air bursts or streams out of the mouth during articulation), whereas *p*, *ph* contrast with *t*, *th* by “place” of articulation, which is determined by the parts of the speech organ that are involved (bilabial vs. dental). A phoneme, Jakobson and Trubetzkoy said, is best represented as a cluster of abstract, hierarchically ordered contrast values, or *distinctive features* (e.g., *v*=+voiced, +fricative, +bilabial; *k*=−voiced, −fricative, +velar). Jakobson later moved to Harvard, and he and his students tried to find evidence from aphasia regarding the phonetic feature hierarchy. Their idea was that the higher a feature on the hierarchy is, the earlier it would tend to break down in aphasia.⁷ They were not particularly successful in finding evidence for this idea.

Some eighty years after Jakobson and Trubetzkoy published their theory, UC San Francisco neurosurgeon Edward Chang and his team used direct cortical recording in the context of brain surgery to demonstrate compellingly

that these very same abstract phonetic qualities could be directly read off the brain. This remarkable experiment involved neurosurgical patients whose brain had been equipped with a matrix of surface electrodes, placed on their auditory cortex for clinical purposes. They volunteered to participate in this experiment, in which they listened to a carefully designed human speech recording, which contained the complete phonetic inventory of English (all its consonants and vowels in different contexts). The electrodes placed on the auditory cortex of their language-dominant hemisphere directly recorded electrophysiological activity during speech perception. Since the brain produces a noisy signal, it took highly sophisticated data analysis by Nima Mesgarani to reveal a remarkable fact: certain clusters of auditory neurons respond not to the sound of a phoneme but rather to its *abstract linguistic features*; that is, speech perception involves immediate abstract linguistic analysis by our auditory cortex. Moreover, it does so along the lines of Jakobson's and Trubetzkoy's theory from the 1930s. Chang and his team read the theory directly off the patients' brains! This remarkable set of results about the linguistic neural pieces, combined with fine neuroanatomy, told a fairly clear story about the abstract nature of phonetic analysis in the human brain and its precise cerebral location.⁸ This is an amazing, neurologically grounded advancement, which has since inspired many attempts to obtain clues regarding the mechanisms by which these neurons carry out abstract featural analysis.⁹ If mapped anatomically, these discoveries would constitute a major step towards a neurolinguistics of the phonetic side of language.¹⁰

Functional Modules: Three Reasons Why Broca's Area Is a Syntax Specialist

Similar discoveries have been made about the brain bases of syntax, a topic of intense investigation and debate since the 1970s. Focusing on receptive language, I will tell you about (A) selective syntactic impairments subsequent to focal brain damage (aphasia); (B) syntactic processes as revealed through brain imaging in health (functional magnetic resonance imaging—fMRI); and (C) direct electrical stimulation of cortical tissue and concomitant syntax testing during awake brain surgery. Results from all three research areas (in which I have been personally involved for decades) indicate that our brain is divided into syntactic areas of specialty.

A. *The Syntax Specialist: Errors in Poststroke Aphasia*

Although expressive language impairment is apparent in individuals with Broca's aphasia, as we could sense clearly through Gabby Giffords's language patterns, testing comprehension can be more easily controlled than other language modalities, such as production, repetition, and naming. As such, comprehension has been studied extensively, with a broad range of sentence types. Controlled experimentation has revealed that the syntax of individuals with Broca's aphasia is partially impaired, evident through investigations of selective error patterns. Most striking is the fine grammatical nature of their deficit, indicating that the brain is structured by syntactic modules: in language comprehension, patients whose Broca's area has been damaged lose the ability to process grammatical transformations properly and do not lose much else.¹¹ Many experiments ask patients to simply indicate who did what to whom in a sentence. Even in such a simple test paradigm, a complex error pattern is observed, displaying a highly selective deficit. For example, patients interpret subject questions, like 1, correctly but miscomprehend object questions, like 2. When asked to answer the question by pointing to one of two boys in a picture, one who the girl is pushing, and the other who pushes her, the patient guesses.¹²

- | | |
|--|-----------------------------|
| 1. Which boy pushed the girl? | <i>Correct performance</i> |
| 2. Which boy did the girl push which boy ? | <i>Guessing performance</i> |
- 

These and related results, coming from patients who are native speakers of languages as disparate as German, Hebrew, English, and Japanese, have supported the notion of a partial syntactic impairment due to a focal, stroke-induced lesion to Broca's area. They have provided important preliminary hints that neurolinguistics must incorporate grammatical principles, as these underlie the brain's modular organization for language, where certain areas specialize in highly specific linguistic functions.

Investigations of stroke victims stop short of revealing a precise anatomical basis for language. Patients with varied lesion types in different anatomical locations have been tested, only exposing a murky anatomical picture of the brain bases for language. The advent of functional neuroimaging gave new hope for more precise human functional anatomy. Imaging technologies can monitor the neural reflexes of cognitive effort in relatively small parts of the brain, which opened the way to noninvasive measurements of neural activity due to linguistic analysis (mainly in the brain's gray matter)

in healthy participants. If pathological language subsequent to brain damage was at center stage until the 1990s, neurolinguistics has since gravitated toward the study of language in the intact brain, mostly through fMRI.

B. The Syntax Specialist: fMRI in Health

Spoken language comprehension necessitates all the stages in language processing: auditory, phonetic, phonological, morphological, syntactic, and semantic (in reading, visual and orthographic analyses may replace some of these). Each step taxes the brain regions that support it. The more complex the analysis, the more taxing, which is what imaging instrumentation attempts to measure. If a damaged area leads to a specific error pattern in aphasia, then in health, this same area is expected to exhibit related activation patterns. Indeed, since the mid-1990s, clues from neuroimaging studies, about the linguistic role of each brain part, have lent support to the view from aphasia and gone beyond it.¹³


A prominent early example of this effort involving the use of fMRI comes from a Max Planck Institute in Leipzig, led by neurolinguist Angela Friederici. Her team tested healthy participants' comprehension of German sentences that contained the very same words as each other but in different orders while conveying (roughly) the same meaning. German syntax allows flexibility in such sentences, and a clever neuroimaging experiment used this property to show that the more MOVE operations a sentence contains, the stronger the neural activity is in Broca's area:

3. a. Heute hat der Opa dem Jungen den Lutscher geschenkt.


Today, Grandpa the child the lollipop gave.

(Today, Grandpa gave the child the lollipop.)

- b. Heute hat ~~dem Jungen~~ der Opa ~~dem Jungen~~ den Lutscher geschenkt.



- c. Heute hat ~~dem Jungen den Lutscher~~ der Opa ~~dem Jungen den Lutscher~~ geschenkt.




These sentences contain zero, one, and two instances of MOVE (annotated by the numbered arrows, with ~~strikethroughs~~ for the original locations of the words). The MOVE complexity increment—from 3a to 3c—determined the incremental activation measured in Broca's area ($3c > 3b > 3a$) and in no other brain region.

Later, neurolinguist Andrea Santi (then at my McGill laboratory in Montreal) conducted sophisticated neuroimaging experiments indicating that the increased neural activity in Broca's area reflects processing difficulty that is not merely due to the linear distance between the origin and destination of MOVE; rather, it turns out that the tissue specifically supports MOVE.¹⁴ In one study, she showed that complexity is not determined by whether the surface order of elements is "canonical" or more frequent in the language. To this end, she kept the length of her sentences, as in 4, constant but manipulated the distance of MOVE (marked by a line) by systematically adding words between the origin and the final destination of the constituent to which MOVE applied; the expectation was that increased distance would increase the perceptual complexity of a sentence, which would lead to increased brain signal in the region that supports MOVE. She then created a set of control conditions by performing the same distance manipulation on sentences with a non-MOVE dependency relation, as in 5: she varied the length of the dependency between a reflexive pronoun *him/herself* and its antecedent noun (marked by a dashed line) by adding words between them, which increased distance.


4. a. The mailman and the mother of Jim loves *the*

woman who Kate pinched *the woman*.



- b. The mother of Jim loves *the*

woman who the mailman and Kate pinched *the woman*.



5. a. The sister of Kim assumes that

the mailman who loves Anne pinched *himself*.



- b. Anne assumes that

the mailman who loves the sister of Kim pinched *himself*.



The distance manipulation seemed a good test of MOVE selectivity in the brain: Santi reasoned that an area specializing in MOVE would activate more strongly to the added complexity in 4b versus 4a, because MOVE, which is presumably computed in Broca's area, would have higher demands in 4b (due to the greater distance between the *bold* and *struckthrough* elements). But the increased distance between the bold elements in the non-MOVE dependency in 5b versus 5a would have no impact on activation levels in Broca's area, because this dependency is not computed in this brain part. The prediction was borne out, as a highly specific MOVE-induced activation was measured in a part of Broca's area and nowhere else in the brain. In a related study, Santi pitted MOVE complexity against embedding complexity. Once again, MOVE but not embedding activated Broca's area.¹⁵

The results from fMRI in health converge on those from aphasia in brain disease: those grammatical operations that are impaired by damage to Broca's area are the very same ones that activate it in health. These results reveal the functional specificity of this brain area, and the tight connection between syntactic MOVE and a specific cluster of neurons.¹⁶

C. The Syntax Specialist: Intraoperative Stimulation of Cortical Tissue

Further corroboration of this approach comes from a recent intraoperative study that tested the MOVE contrast with direct electrical stimulation mapping during awake brain surgery.¹⁷ As the brain has no sensory innervation (it doesn't sense itself), patients can be awake while their brain is operated on, and intraoperative cognitive testing is sometimes surgically required. Direct electrical stimulation, used for clinical purposes, temporarily blocks a small brain area, while the awake patient is presented with a task. In this study, when Broca's area was stimulated, the comprehension of sentences with MOVE, like 6b, but not of their controls, like 6a, was transiently disrupted in several patients—they made a large number of errors on the passive sentence but hardly any on its active counterpart.

6. a. The boy is kissing the girl.
 b. *The boy* is kissed ~~the boy~~ by the girl.



Importantly, when the same stimulation procedure was repeated in other cortical loci (most notably, the left temporal lobe, which is roughly where another language region—Wernicke's area—is located), the same transient blocking did not lead to a language disruption. This highly specific error behavior lends further support to the modular, MOVE-based approach to Broca's region highlighted above.

Beyond Syntax

The fine distinctions just reported were later augmented by many others, including a connection between areas adjacent to Broca's area and highly specific and complex semantic functions.¹⁸ For example, work in my laboratory recently uncovered the neural basis of a central logicosemantic function: negation (a.k.a. inference reversal through the insertion of words like *no*, *not*, *less*, *few* to a sentence). Tests we carried out in fMRI found that the

contrast between sentences with and without negation selectively activates a small area in the left anterior part of the insula—a cortical structure in the frontal lobe that is somehow “buried” deeper than most cortex.¹⁹

Small Anatomical Modules for Language

These are mere pieces of phonetics, syntax, and semantics that constitute small functional modules. But what about their neural substrate? Can related anatomical pieces or modules be found and defined in the human brain? Can Broca’s area (or parts thereof) be localized precisely? Does a function reside in a single locus or in many? A quick tour into the world of microscopic anatomy might help.

Functional anatomy requires precision. Humans feature large individual variation in brain structure (our noses, chins and foreheads are shaped differently, so why expect our brains to look alike?). To get around this variation and to provide a characterization of the brain’s functional anatomy that is precise from both the function and anatomy perspectives, a method is needed that divides, or *parcellates*, the brain into coherent, function-sensitive pieces and that at the same time transcends interpersonal variability. Advances in computational neuroanatomy turn out to be very helpful here.

Our brain is famous for its complex visible surface terrain, whose convoluted topography has not been useful in dividing it into functional pieces (for one thing, about two thirds of the cortical tissue is buried in the folds, or sulci; it can only be viewed through dissection or imaged by algorithms that transform the brain into an inflated object with smooth surfaces, which entails the loss of some three-dimensional information). But the brain’s cortical tissue also has rich microscopic structure, which has been explored since the late nineteenth century, an effort pioneered by Nobel laureates Santiago Ramón y Cajal and Camillo Golgi, whose work, mentioned in chapter 3, enabled the selective staining of densely organized neural tissue.²⁰ In the human context, Oskar and Cécile Vogt and their student Korbinian Brodmann partitioned gray matter (cortex) from white matter (fiber) and, based on the arrangement of cell layers in cortical tissue and the different cell shapes and types they contain, they parcellated the brain’s cortex, or gray matter, into a mosaic of well-delineated areas. In

other words, they used microscopically visible patterns that cortical tissue exhibits to build a cortical mosaic (a.k.a. “cytoarchitecture”).²¹

A cytoarchitectonic area reflects local structural uniformity, which distinguishes it from its neighbors. The uniformity of a small region is believed to reflect a coherent network, from which a set of specific functions emanates. These areas are therefore the anatomical pieces that we expect to be congruent with functional components, in our case, pieces of cognitive function. Indeed, the functional relevance of the cytoarchitectonic mosaic became evident when sensorimotor and visual maps in mammals were constructed. Later, the borders of this mosaic correlated with the border of a strong marker of function—neurotransmitter receptor distribution within the cortical ribbon. A 1909 human cytoarchitectonic atlas by Brodmann consisted of 43 numbered regions (Brodmann areas—BAs). Its present-day descendants use related—though more detailed—nomenclature.²²

The modern remapping of Broca’s area (BA 44 and BA 45) by Katrin Amunts and the late Karl Zilles, of the University of Düsseldorf’s Cécile and Oskar Vogt Institute and the Jülich Research Center in Germany, marked a new era in microscopic mapping of language and other abilities. It has since become part of a high resolution atlas of the human brain that divides cortical tissue into dozens of small modules, delineated through the use of microanatomical criteria. Figure 7.4 shows two adjacent cortical sections. Layers are marked (though not always visible to a nonprofessional eye), and the different cellular arrangements are apparent. Computational methods determined the boundary between the areas.²³

Figure 7.5 displays a three-dimensional representation of the surface division of human cortex into areas in the Jülich cytoarchitectonic atlas of the human brain. Highlighted in red and blue are the parts BA 44 and BA 45, from which the slices in figure 7.4 were taken. This atlas is the best currently available method for cortical parcellation and subsequent localization of functional modules within anatomical ones.²⁴

How good is the match between the anatomical and functional pieces? After all, this is at the heart of our research program, a first step in a principled relation between anatomy and function. The match can be measured by using the anatomical parcellation maps as a background to the functional localizers. To conclude that an anatomical piece is the “seat” of a functional one, we need good overlap between the anatomical (cytoarchitectonic)




Figure 7.4

Sections from the two cytoarchitectonic parts of Broca's area, BA 44, 45. Laminar structures are different. Courtesy of Katrin Amunts, Jülich Research Center.




Figure 7.5

Parcellation of human cortex by cytoarchitectonic borders. BA 44, 45 are highlighted. Courtesy of Katrin Amunts, Jülich Research Center.

areas and the clusters that mark functional activation in fMRI experiments. We extract those pieces of the brain's anatomy that are known for their involvement in syntactic analysis; these are shown in figure 7.6A (red and yellow blobs). Next, we return to Santi's syntax study, which carefully isolated the brain area that supports MOVE, not other syntactic dependency relations, as well as the semantic study in my lab that isolated negation. We first asked whether Santi's activated area overlaps with a coherent piece of anatomy. We thus created a joint map, a slice of which is shown in figure 7.6B, with colored blobs from two different sources: the anatomical area, taken from 7.6A (red blob in 7.6B), and the area that was modulated by the function, drawn from Santi's fMRI results (green blob).

The image reflects the clear match between the two. The cluster of MOVE-activated neurons belongs to a single microanatomical area. The visual impression is confirmed computationally. Similar results are obtained in Santi's study on complexity, as well as in Friederici's earlier study. All three implemented the same MOVE contrast through different materials, task, controls, and language, and BA 45 kept recurring as the one area specializing in MOVE.

The negation map, discovered in my laboratory, localized to a small cortical area adjacent to Broca's area (blue blob in figure 7.6A). Again, the neural underpinnings of negation were detected, through joint work with anatomist Katrin Amunts and her Jülich team. The functional area shown in figure 7.6C was congruent with a cytoarchitectonically coherent part of the anterior insula of the left hemisphere, shown in figure 7.6D.²⁵ This finding provided yet another instance of the match between highly specific cognitive functions and well-defined small anatomical areas. A microscopically defined map of linguistic functions is beginning to emerge.

Is this spatial adjacency between syntactic and semantic functions accidental? I believe it is not. The "negation area" is located midway between Broca's and the anterior cingulate gyrus, where reasoning processes are known to take place.²⁶ As reasoning is often (though not always) done over abstract representations of linguistic content, the result suggests that the left anterior insula is a medium for these representations, which are then fed forward to the cingulate. On this interpretation of the results, the left anterior insula receives the output from Broca's area, organizes it in a logical format (i.e., like a logical formula), and it is then sent to the anterior cingulate, where reasoning operations take place. Such pathways are critical for




Figure 7.6


Hints regarding the microanatomy of syntax and semantics, represented in two-dimensional brain slices and a three-dimensional rendition. **A.** BA 44 (red), 45 (yellow), which together comprise Broca's area, and Id7 (blue), part of the anterior insula. All blobs are located in the left cerebral hemisphere. **B.** A typical fMRI activation cluster marking areas that support MOVE (green) superimposed on BA 45. **C.** The insular brain area that has been found to support negation. **D.** A three-dimensional depiction of the functional area in C (magenta), superimposed on the anatomically delineated area Id7. The dots represent the centers of mass of the anatomical and functional blobs, showing that they are nearly identical (difference in blob size is merely due to a "thresholding" technique). Reproduced by permission from A. Santi and Y. Grodzinsky, "Working Memory and Syntax Interact in Broca's Area," *NeuroImage* 37, no. 1 (2007): 8–17; Y. Grodzinsky, I. Deschamps, P. Pieperhoff, F. Iannilli, G. Agmon, Y. Loewenstein, and K. Amunts, "Logical Negation Mapped onto the Brain," *Brain Structure and Function* 225 (2020): 19–31.

an understanding of the neural basis of semantics and logic. Any account of the computational structure of language must include this pathway.

In sum, when we examine the current state of evidence, the neural bases for language appear to be organized as a mosaic of relatively small but connected clusters of neurons. Extant empirical findings (of which I selected just a few, simply because they associate microanatomy with microfunction) identify a clear functional-anatomical unit, one of several that have thus far been uncovered. It should be said that “small” here still translates to a large number of neurons: a quick calculation based on current data indicates that the average number of neurons we have in each of BA 44 and 45 is about 3×10^8 (three hundred million) neurons.

The approach I have described, while highly structured and logically transparent, has faced a series of objections (a putative LLM-based alternative is presented in the next chapter). Critics have claimed that regularities that seem to be observed are simply not there; that is, that the brain is more variable than depicted here.²⁷ The response to these objections notes that the studies they cite used linguistically unsystematic tests, which simply blurred the (otherwise clear) picture. Similar criticisms have been based on reviews of fMRI studies that claimed to find much variability in the localization of syntactic processes.²⁸ Are variable results due to nonuniform choice of test materials or to natural interparticipant variability?

A decisive answer seems to come from a systematic review of fMRI studies that my colleagues and I recently conducted on the neural bases of the syntactic MOVE component.²⁹ We scrutinized multiple fMRI studies of syntax, and selected those experiments that chose properly controlled test materials that focused on the MOVE component of sentence processing. This selection gave a clear answer, based on the brain activity detected in over 250 participants: as figure 7.7 shows, when a MOVE contrast was deployed experimentally, left Broca’s area (BA 44, 45, marked red and orange) was the only brain area activated in every single study (activation points, or maxima, are superimposed on the cytoarchitecture). This was true of a variety of MOVE contrasts in English, German, and Japanese. Despite noise and false starts, then, the conjecture that specific pieces of syntax reside in designated brain pieces appears vindicated by investigations in a broad range of experiments.

**Figure 7.7**

Aggregated results from 13 experiments with the MOVE contrast, superimposed on a partial cytoarchitectonic color map (see legend on right). Pins mark the maxima of MOVE activation clusters, colored by language (English=orange, Japanese=green, German=magenta). Left Broca's area (yellow and orange) was the only brain region activated in all studies. Other activation points were scattered, some in and around Wernicke's area (purple, blue and green). Such activations, however, were documented in less than half of the studies.

A Divide-and-Conquer Program and Agenda for Computational Neurolinguistics

The approach to brain-language relations that combines anatomy and linguistic theory appears vindicated. Indeed, the neuroscience of language has come a long way since Broca, Wernicke, and Geschwind: their notion of pieces of cognition in brain pieces was right, but they had no precise anatomical methods, and they missed the fact that the brain is a linguist—that its microanatomy is shaped (though perhaps not exclusively) in a manner that honors linguistic structure. Indeed, we can conclude that linguistic and microanatomical tools are crucial for the right story of the talking brain. Linguistic tools parse the function into its component parts, while

cytoarchitecture parcellates cortex into functionally meaningful anatomical entities. The extant state of knowledge admittedly stops short of delivering actual neuronal models that turn structure into language functions. It lacks a description of the connectivity between areas, and it suffers technological limitations. Precise functional brain mapping in humans, especially if based on microscopic cytoarchitectonic considerations, is difficult to achieve in vivo, and can only be obtained via postmortem dissection. As functional maps require a living brain, the sampling of each individual participant for both anatomy and function is therefore not an option. Pending further technological innovations, we have to live with the fact that the resolution currently goes as high as figure 7.7 allows. Thus progress may be slow and halting, but this direction is likely the right one. Limited resolution need not limit our vision, as it did not limit early brain mappers such as McGill's neurosurgeon Wilder Penfield, who in the 1930s pioneered the mapping of the human sensorimotor cortex through the use of intraoperative direct electrical stimulation. The resulting representations, better known as *homunculi* (little persons), indicate that bodily organs are (disproportionately) represented in distinct but adjacent brain loci in both sensory and motor cortices, in a manner that later turned out to have fairly clear geometric properties.³⁰ Visual and auditory maps also have such structure (known as *retinotopy* and *tonotopy*) in humans as well as lower species, and we have seen precise functional (though not anatomical) maps of phonetics.

The Modularity of Language: Cytoarchitectonic Areas as a Mosaic of Linguistic Homunculi

The results presented here point to a fine mosaic that is the functional anatomy of language, made up of small, stable, and cytoarchitectonically coherent cortical substrates for major sentence processing components, discoverable only upon controlled linguistic manipulation. These small pieces, whose microanatomical makeup is known, should give rise to a detailed cortical map of linguistic mechanisms that can be modeled computationally, by anatomy-guided local networks.³¹

This program is very much in keeping with ideas that I first heard many years ago: when I was a student, a young fellow named Geoffrey Hinton came to the Center for Cognitive Science at MIT, to deliver a talk about his work on networks and parallel computation—the work for which he has just been

awarded the Nobel Prize in Physics. Sitting in the audience were celebrated philosophers Daniel Dennett and his friend Jerry Fodor, who had just published his programmatic monograph *The Modularity of Mind*, in which he espoused a view of the cognitive system that divided it into functional modules, language and its components being a central one among these.³² At the end of Hinton's talk, Fodor got up and asked (in his very special, somewhat intimidating, tone): "So, what about modularity?" Blushing, Hinton replied that his network perspective did not contradict the idea behind *The Modularity of Mind*. As far as he could tell, Hinton said, the human mind may well be composed of small modules, whose internal makeup is a set of connected networks of the kind he had presented. Fodor then said "Good!" and sat down.

I am not sure that Hinton would endorse his early position today, but his reply echoed in my mind for decades, during which I was wondering about ways to implement the modular view of the language system and of the brain. I now realize that this view is not a fantasy. The most plausible view of the linguistic brain must be that it consists of local networks that constitute small and identifiable anatomical homunculi, dictated by microscopic anatomy. These local networks power highly constrained linguistic principles and rules. As I will describe at the beginning of the next chapter, this is exactly the current practice in neurologically realistic network models of vision. The path may not be easy (and LLM investigators do not currently go there), but the evidence suggests that it is there, for us to discover.

Taking Stock

This chapter described the linguistic perspective on brain–language relations. We saw:

- The tragic yet courageous story of Gabby Giffords, who survived an assassination attempt but remained with Broca's aphasia
- Broca's, Wernicke's, and Geschwind's early language localization maps. Language is divided by the modalities through which it is practiced (speaking, listening, reading, writing)
- Distinctive phonetic features and their brain localization
- That the syntactic rule `MOVE` is well-localized in Broca's area, as seen in Broca's aphasia, in fMRI studies, and even through a test of `MOVE` in awake surgery. Other parts of syntax are located elsewhere
- Hints about semantic localization from the neuroscience of negation
- That anatomical parcellation of cortical tissue is done on the basis of microscopic definitions of cytoarchitectonic borders
- That pieces of linguistic theory and anatomical pieces are congruent in brain pieces that may contain up to 3×10^8 neurons each
- The future: more detailed neurolinguistics, better high resolution anatomy
- A modular network-based program for neurolinguistics: the quest for linguistic homunculi, akin to mapping projects in the sensorimotor systems, as well as vision and audition

This is a section of [doi:10.7551/mitpress/15932.001.0001](https://doi.org/10.7551/mitpress/15932.001.0001)

How Deeply Human Is Language?

Chomsky, the Brain, and the AI Fantasy

By: Yosef Grodzinsky

Citation:

How Deeply Human Is Language?: Chomsky, the Brain, and the AI Fantasy

By: Yosef Grodzinsky

DOI: 10.7551/mitpress/15932.001.0001

ISBN (electronic): 9780262052023

Publisher: The MIT Press

Published: 2026



The MIT Press

The MIT Press
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
mitpress.mit.edu

© 2026 Massachusetts Institute of Technology

All rights reserved. No part of this book may be used to train artificial intelligence systems or reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data is available.

ISBN: 978-0-262-05200-9

EU Authorised Representative: Easy Access System Europe, Mustamäe tee 50, 10621 Tallinn, Estonia | Email: gpsr.requests@easproject.com