Published in final edited form as:

Cogn Neuropsychol. 2012 February ; 29(1-2): 34–55. doi:10.1080/02643294.2012.710600.

The *maps problem* and the *mapping problem*: Two challenges for a cognitive neuroscience of speech and language

David Poeppel

Department of Psychology, NYU, 6 Washington Place, New York, NY 10003, 212 992 7489,
david.poeppel@nyu.edu

Abstract

Research on the brain basis of speech and language faces theoretical and empirical challenges. The majority of current research, dominated by imaging, deficit-lesion, and electrophysiological techniques, seeks to identify regions that underpin aspects of language processing such as phonology, syntax, or semantics. The emphasis lies on localization and spatial characterization of function. The first part of the paper deals with a practical challenge that arises in the context of such a research program. This *maps problem* concerns the extent to which spatial information and localization can satisfy the explanatory needs for perception and cognition. Several areas of investigation exemplify how the neural basis of speech and language is discussed in those terms (regions, streams, hemispheres, networks). The second part of the paper turns to a more troublesome challenge, namely how to formulate the formal links between neurobiology and cognition. This principled problem thus addresses the relation between the primitives of cognition (here speech, language) and neurobiology. Dealing with this *mapping problem* invites the development of linking hypotheses between the domains. The cognitive sciences provide granular, theoretically motivated claims about the structure of various domains (the ‘cognome’); neurobiology, similarly, provides a list of the available neural structures. However, *explanatory* connections will require crafting computationally explicit linking hypotheses at the right level of abstraction. For both the practical *maps problem* and the principled *mapping problem*, developmental approaches and evidence can play a central role in the resolution.

Keywords

localization; computation; cognitive neuroscience; linking hypothesis

Introduction: A practical-technical challenge and a principled-conceptual challenge

Research on the neurobiological foundations of cognition, in general, and speech and language processing, in particular, faces a variety of interesting empirical and theoretical challenges. Two problems are discussed here, a practical one and a principled one. The practical problem has to do with how we should conceive of (one of) the main forms of data that lie at the basis of cognitive neuroscience: maps of the brain and maps of brain activation. This *maps problem* concerns the extent to which spatial information about brain activity provides satisfactory descriptions of the neural basis of perception and cognition. The techniques that currently dominate the field (whether spatially specialized, such as fMRI, or temporally specialized, such as MEG) predominantly characterize data in terms of spatial attributes (such as local topographic organization like retinotopy or somatotopy, processing streams like dorsal versus ventral pathways, or networks of interconnected brain regions). Thinking and talking about brain activity in spatial terms is very intuitive (spot A ‘does’/houses/executes function X, spot B does Y, etc.), often more or less correct (at some

level of description), and has captured both the professional the popular imaginations. Visually compelling publications such as *Images of Mind* (Posner & Raichle 1997), *Mapping the Mind* (Carter & Frith 2000), or *Portraits of the Mind* (Schoonover 2010) elegantly summarize the successes of the research program. They also clearly reflect (in their titles as well as their substantive content) the deeply held, implicit presupposition that making a visual, spatial map of mental faculties is a critical step in a proper explanatory account of the brain basis of psychological function. This stance is likely to be correct *in part*, but at best will reflect an incomplete understanding; uncontroversially, localization and spatial mapping are not explanation (e.g. Poeppel 2008). In a typical cognitive neuroscience study of speech, lexical, or sentence-level processing, participants will engage in some naturalistic or artificial task while their brain activity is monitored. The analyses show that some area or areas are selectively modulated, and it is then argued that activation of a given area underpins, say, phonological processing or lexical access or syntax. Although this superficial characterization underestimates the state of the field, it is fair to say that the canonical results – very much at the center of current research – are correlational; that is to say, systematic relations consistently occur between brain areas and some functions that reappear across studies, but we have no *explanation*, no sense in which properties of neuronal circuits that we understand account for the execution of function. How to proceed? I suggest that even high-resolution data from (existing or to be developed) new techniques will remain inadequate unless we decompose the cognitive tasks under investigation into computational primitives that can be related to local brain structure and function, in a sense instrumentalizing the computational theory of mind (e.g. Pinker 1997) more aggressively. We will need to seek theoretically well-motivated, computationally explicit, and biologically realistic characterizations of function to take the work to the next level. Some of the issues that arise this central part of cognitive neuroscience, making maps, are discussed in the next section, with a view towards how maps implicate functions, and a view towards how a developmental perspective can sharpen the issues and help construct productive links between structural description and functional analysis.

The principled problem deals with the ‘alignment’ between the putative primitives of cognition and neurobiology and constitutes a more abstract challenge. Addressing this *mapping problem* – what is the relationship between the ‘parts list’ of cognition and the ‘parts list’ of neurobiology? – is considerably more difficult than it might seem at first, ultimately requiring the development of appropriate linking hypotheses between the different domains of study. I use the expression *mapping* here not to refer to the assignment of putative linguistic or psychological function to brain areas, be they distributed or localized, microscopic or macroscopic. Rather, I mean with mapping the investigation of the (ultimately necessary) *formal relations* between two sets of hypothesized inventories, the inventory constructed by the language sciences and that constructed by the neurosciences: how do the primitive units of analysis of the cognitive sciences map on to the primitive units of analysis of the neurosciences? The cognitive sciences, including linguistics and psychology, provide detailed analyses of the ontological structure of various domains (call this the ‘human cognome,’ i.e. the comprehensive list of elementary mental representations and operations), and neurobiology provides a growing list of the available neural structures. To exemplify, the infrastructure of linguistics – building on formally specified concepts such as *syllable* or *noun phrase* or *discourse representation*, etc. – provides a structured body of concepts that permit linguists and psychologists to make a wide range of precise generalizations about knowledge of language that speakers bring to bear, about language acquisition, online language processing, historical change, and so on. Similarly, the infrastructure of the neurosciences – drawing on units of analysis such as *dendrite* or *cortical column* or *long-term potentiation* – captures a variety of structural and functional features of the brain, with profound consequences for the neural basis of cognition and perception. However, how do the hypothesized units of analysis relate? Simple reductionist assumptions

cannot even be stated in any sensible way (say, e.g., ‘neuron = syllable’). Such alignments may seem potentially pleasing at first glance - but that pleasure subsides quickly if one has to actually develop such a view and do the work of representation and computation. The fact of the matter is that we have very little to no idea how the stuff of thought relates to the stuff of brains, in the case of speech and language – and virtually all other cases (Gallistel & King 2011; Mausfeld 2012). In section 3, I argue that in order to link cognition and neurobiology in an explanatory fashion will require the formulation of computationally explicit linking hypotheses at the right level of abstraction. With respect to this principled, issue, too, developmental data provide critical insights, for example in specifying the parts lists attributed to the infant from the start (innate, precompiled toolbox) versus the representations and operations argued to be acquired and constructed.

The guiding perspective throughout is that of a Marr-style approach (Marr, 1982). Separating computational, algorithmic, and computational levels of analysis, rather than being anachronistic, will be helpful in better defining the nature of the challenges and pointing to some tentative answers, or at least plausible approaches to possible answers. Cognitive neuroscience has made immense progress in the last 20 years, and the grounds for optimism are justified. However, what constitutes an explanation is rarely addressed. Adapting the position articulated by Marr (and others) on the distinct but tightly yoked levels of description, I suggest that a focus on the algorithmic and representational level can provide a potentially productive perspective in seeking hypotheses that bridge between high-level computational and low-level implementational characterizations.

In the next section, on maps, I outline four types of spatial organization that form the basis for much of contemporary research on the brain basis of speech and language: local regional organization of a given brain region (e.g. ‘Broca’s region’), the notion of processing streams (e.g. ‘dorsal’ versus ‘ventral’ streams), hemispheric asymmetries, and putative networks of connected brain areas. This ‘brain mapping’ program of research, emphasizing different ‘functional anatomic units of analysis’ ranging from the microscopic to the network level, has been productive and successful and remarkably popular in the public understanding of cognitive neuroscience - but it needs to be considered what the next steps could be to cash out the promise of specifying such spatial maps. After discussing the issues surrounding the maps problem and suggesting some questions for which developmental evidence can and should play a critical role for progress, I turn to the mapping (or alignment) problem in section 3, developing the argument why it is so difficult to establish the principled relations between linguistic or psychological primitives and neurobiological primitives. The paper closes (section 4) with some tentative suggestions about how one might proceed in a way that could bridge the *maps* and *mapping* problems, endorsing a computational approach that attempts to identify a level of analysis that might facilitate at least the formulation of linking hypotheses.

The standard research strategy in cognitive neuroscience: making maps

Defining the spatial layout of brain organization has a long and successful history. In fact, perhaps the first large-scale model of brain organization in systems neuroscience was the so-called ‘connectionist model’ for language based on the seminal insights by Broca (1861) and Wernicke (1874), arguing that there are two localized centers that underpin language production and comprehension (Geschwind 1970). (Though this model has been unbelievably influential to this day, especially in clinical contexts, it is widely accepted that it remains woefully underspecified, both from the linguistic and neurobiological points of view. For discussion, see e.g. Poeppel & Hickok 2004; Grodzinsky & Amunts 2006.) Spatial organization of one form or another has been investigated profitably at many levels of description, from determining (i) at a microscopic level the features that specify local neural

analysis within a brain region (say, e.g., retinotopic mapping in the visual system) to (ii) information processing streams (e.g. the *what* versus *where* systems in vision) to (iii) sophisticated network analyses (Sporns 2012) that model how distributed regions are connected in support of function and (iv) attempts to understand the human connectome (Seung 2012), that is, the comprehensive mapping of how all areas are connected, cell by cell: the total wiring diagram of the human brain.

Why does the search for finding a map with local topographic organization induce such yearning in the field? What underlies the seductive appeal of map-making? Maps of one form or another are, of course, on target for sensory and motor systems, revealing fundamental organizational principles that have helped model aspects of perception and motor control. Searching for topographies has been extremely successful in the case of vision (e.g. retinotopic organization of numerous visual areas), in the case of the somatosensory system (the sensory homunculus), widely attested in auditory processing (tonotopic mapping), and so on. So, in a strategic sense, aiming to understand the spatial organization of *cognitive* functions builds on well-established neuroscientific observations.

The study of the visual system has formed the conceptual basis for many aspects of the standard paradigm in cognitive neuroscience. Much of the original work is elegantly summarized in a monograph by David Hubel, one of the chief architects of visual neuroscience *Eye, Brain, and Vision* (Hubel 1995). In that domain, it was assumed that the external visual world is represented by decomposition into elementary features (or basis functions) and subsequent composition into higher-order, complex visual representations. (The model of vision espoused is ultimately underspecified and no longer the dominant view; however, it has been a remarkable productive empiricist paradigm.) Local features (e.g. retinal position, edges, luminance) are aggregated along feed-forward pathways, leading to the construction of more complex visual representations (e.g. faces, as an obvious example). Three generalizations emerged from this work. First, decades of visual neurophysiology confirm that there is *localization of function*: neurons and ensembles of neurons perform operations that are regionally specific. Cells in an early visual area such as primary visual cortex (V1) have firing properties that reflect strikingly different receptive field size and response selectivity than cells in a higher order area, say in the infero-temporal cortex. The operations underlying visual decomposition (and composition) are localized and specific. Second, the feed-forward pathways along which the visual information is aggregated suggest *hierarchical processing*. Neurons along a processing stream reflect ever-more integrated (in space, time) and abstract properties of the external world. Third, there exist *concurrent pathways* for information processing that underlie different larger-scale functions. Originally, these were conceived as what versus where streams (Mishkin et al. 1983); subsequently the discussion has focused on object identification versus action, sensorimotor transformation, and other visual tasks (Goodale & Milner 1992). Cumulatively, this research on the neurobiological foundations of vision has profoundly influenced the way of asking questions about the neural basis of perception and cognition.

But it is not at all clear how such spatial thinking plausibly extends to cognition, and especially language. Research on the brain basis of language has (implicitly) adopted the approach of vision, going as it were ‘from the outside in,’ assuming that sounds (or signs) are analyzed in one region, aggregated into morphemes in another, followed by some sequential syntactic processing in yet another region, and so on. In fact, an enormous part of the neuroscience of language literature is dedicated to functional localization at a very high level, seeking to identify segregated regions for speech perception, phonology, morphology, syntax, semantics, etc. But it is an empirical question whether *any* spatial mapping principles as they appear in sensory systems extend to cognition or language. (Verbs here, nouns there? Phonemeotopy? Spatial distribution of syntactic categories? It is not even clear

if such questions can be properly formulated in a theoretically motivated context – although precisely these examples have been discussed in the cognitive neuroscience literature; Pulvermueller et al. 1999, Obleser et al. 2003, Hillis & Caramazza 1995) We have, of course, learned a great deal about the likely functional contribution of certain brain regions (e.g. the superior temporal gyrus), processing streams (e.g. dorsal and ventral streams in speech), hemispherically organized brain structures, and so on. The examples of spatial organization raised in the next sections are intended to be forward looking, hoping to extend what spatial maps of the future might tell us. The point to bear in mind throughout is that the spatial organization of information processing systems can be a useful, and even necessary, intermediate step in explaining a system, but even fantastic localization of function, incorporating the newest techniques (e.g. multi-voxel pattern classification of imaging data or various new approaches to electrophysiology), does not constitute an explanation (Poeppl 2008). The cartographic imperative does not suffice. It is the mechanistic understanding of the *function* that we seek, and that is not going to be tractable by localization of function and spatial topographic mapping alone.

In what follows, I discuss four examples that describe progress in speech and language research in the context of spatial mapping studies and that exemplify the challenge of relating brain maps to language processing. One lesson from the perspective developed here is that a traditional, quasi-modular view (assigning phonology here, syntax there, etc.) cannot succeed, in part because the spatial studies suggest that much more fine-grained computational decomposition of linguistic domains into elementary representations and computations will be required.

Maps of regions

Two examples of brain regions are discussed. First, the canonical language region, Broca's area, is examined. Second, we consider the contribution to language processing of visual areas located early in the processing hierarchy.

Subsequent to the original definition by Broca (the third frontal convolution of the inferior frontal gyrus), the first detailed anatomic analysis of the inferior frontal gyrus was provided by Brodmann (1909). He argued for a division of that region into two or three subregions; typically, the cytoarchitectonically distinct Brodmann areas 44 and 45 are included (some authors also count area 47). Roughly until 2000, generalizations about Broca's area – whether deriving from imaging data or from neuropsychological deficit-lesion correlation data – were made at best at that level of analysis (see Grodzinsky & Amunts, 2006, for an extensive review). That is to say, researchers either made claims about 'Broca's area' - or on occasion about areas 44 versus 45 - in the interpretation of their experimental data. The resolution of the analysis was very rarely higher than that. There exist some interesting connectivity studies (e.g. Anwander et al. 2007), but the functional anatomic characterization has been coarse.

New anatomic techniques (imaging as well as immunocytochemistry) have been applied more recently, and data in an influential paper by Amunts et al. (2010) shows that the organization of this one cortical region is much more complex, incorporating, depending how one counts, 5-10 local areas that are distinguished by their cellular properties. Figure 1, taken from from Amunts et al., illustrates the currently hypothesized organization of Broca's region, highlighting anterior and posterior as well as dorsal and ventral subdivisions of known fields and also pointing to the additional opercular and sulcal tissue that has been implicated.

Now, it is extremely likely that each subregion performs at least *one* different local computation – after all, they are anatomically distinct. And it is worth remembering that we

have not even scratched the surface of possible laminar differences within and across these cortical fields. (Recall that cortex is a sheet of layers, each of which is distinct in terms of cell types, receptor distribution, inputs and outputs, etc. We know from the study of other regions, using animal models, that there are intricate feedforward and feedback projection schemes, canonical cortical circuits (Douglas & Martin 1991), and a rich infrastructure mediating locally distinct organization.) Suppose then, for the sake of argument, that there are merely five anatomic subdomains in Broca's region, and suppose that each anatomic subdomain supports 2 types of operations. These are rather conservative numbers, but we must then still identify 10 types of computations. (This is likely to vastly underestimate the actual complexity.) Some of the computations will be applicable to any input data structure (structured linguistic representation, speech signal, or other), since it is well established that Broca's region is engaged by numerous cognitive tasks (Embick & Poeppel 2006). Other computations may be dedicated to data structures native to the linguistic cognitive system. We know next to nothing about precisely what kind of conditions need to be met to be processed by Broca's region, and in particular by one of the many areas that constitute this part of the inferior frontal cortex.

So how to proceed? If other examples from the central nervous system can serve as a guide (say the retina with its many cell types and associated functions, or primary visual cortex with its ocular dominance columns and orientation pinwheels), Broca's region is likely either (i) to support a number of different types of operations/computations, or (ii) supports similar operations on very different input data structures to yield different outputs. In either case, it is our job to acquire and describe data at a much higher resolution than has been customary to do justice to the known organizational complexity, at least at the level of known anatomic distinctions as shown by Amunts et al. (2010). (Imagine, for comparison, studying visual perception without acknowledging the role primary visual cortex plays in eye-specific and orientation-selective processing; it would be bizarre.) How can this be achieved? In the next decades, we will presumably have access to new techniques with higher spatial resolving power. And existing but relatively new analytic techniques (say multi-voxel pattern classification in the context of fMRI) provide also will provide a way to make necessary anatomic distinctions. In the meantime, we will rely on invasive approaches (e.g. Sahin et al. 2009), animal studies (that are, of course, limited in their value with respect to questions of speech and language), and the coarser techniques that we currently use in cognitive neuroscience. But notwithstanding any technical progress, it is quite clear that generalizations of the type 'Broca's region supports syntax' or 'Broca's region supports hierarchical processing' or 'Broca's region underpins verbal working memory' are idealizations or simplifications that are not just insufficient but possibly misleading. What appears to be required is a much finer decomposition, and the question is what the most appropriate level of decomposition might be. This is an issue discussed in the context of the *mapping problem* below, but the challenges of the *maps problem* already point to the nature of the matter.

In acknowledging the architectural complexity of Broca's region as depicted in Figure 1, it becomes clear that the classical model of brain and language – Wernicke's region (e.g. for comprehension), Broca's region (e.g. for production), and the arcuate fasciculus connecting them – must be abandoned (see Poeppel & Hickok 2004 for discussion). Virtually any angle from which one examines the classical model (linguistic, psycholinguistic, neurobiological, computational) undermines its utility.

Next, let us turn to another brain region, occipital cortex, and a different type of 'local region' issue. The early visual areas are not typically associated with language processing, but recent research reveals how processing linguistic materials implicates very early visual regions and how language experience shapes early visual processing. Contrary to intuition,

the data from neurophysiology suggest that even unquestionably domain general regions such as early visual cortex participate in rather abstract aspects of processing, not just the superficial, feedforward perceptual analysis of input signals. To be clear: it is not surprising that visual areas play a role in processing visually presented language. There exist, of course, relations between linguistic representations and sensory representations in vision (text) – but also in hearing (speech), and touch (Braille). At stake is, rather, whether the visual activations observed in some recent studies can help illuminate the computational subroutines and constitutive representations that are invoked during language processing. The relevant concept that connects the studies is predictive processing.

How does language processing generate neural activity in early visual areas? Unlike Broca's region, early visual areas are decidedly not counted among the typical language areas. Yet recent MEG studies by Dikker, Pykkänen, and colleagues show that syntactic and lexical-semantic cues in the input affect neural activity in occipital cortex between 100-130 ms – that is to say, in very early visual areas (Dikker et al. 2009, 2010; Dikker & Pykkänen 2011). These authors experimented with expected versus unexpected word *categories* (syntactic predictions) in sentential contexts and showed enhanced activity in visual cortex by 120 ms. Similarly, lexical-semantic predictions affected early visual processing, again by 100 ms. They interpret these findings, plausibly, in the context of predictive models of language processing: the brain generates (syntactic, semantic) predictions about upcoming input whenever any prediction is possible; some of these predictions might be actualized as form-based estimates available to sensory areas (say, like an orthographic or even more basic code). In other words, these regions are not language areas but do reflect consequences of the predictive nature of online language processing.

A different type of predictive effect, one driven by linguistic experience and reflecting rather profound developmental plasticity, is demonstrated in electrophysiological results by Almeida et al. (forthcoming). They studied the MEG responses in visual cortex elicited by possible and impossible gestures in signers and non-signers. Both subject populations showed completely canonical responses to faces and inverted faces. However, when viewing signs and gestures, although both groups were able to detect carefully crafted biological anomalies in the gestures, the congenitally deaf signers showed marked changes and facilitation (in terms of sensitivity, reaction time, and neural activation). That work provides new evidence that high-level visual-perceptual properties related to human form processing can be detected within the first 100 milliseconds of visual processing - and that this early processing has functional significance. The data show modulatory effects within the vicinity of primary visual cortex (V1), a neural region that historically has been considered to be reactive to only low-level visual properties. The altered response properties are driven by linguistic experience. The findings derive from comparing the visual processing capabilities of normal hearing listeners with profoundly deaf individuals, whose necessary reliance on visual information makes them plausible candidates for significant cortical reorganization of the visual system. Profoundly deaf signers show exquisite sensitivity to the detection of these gestural manipulations, unlike normally hearing subjects. These data suggest that sophisticated, rather abstract visual constraints may be realized remarkably early in sensory processing streams. It is argued that these constraints may be a reflection of the encoding of *linguistic* properties of signed languages, reflecting a visually driven internal forward model for human gestures that recruits even primary visual areas. As in the other examples of visual cortex activation, the interpretation centers on the predictive nature of the language system, whether it is studied in written English or signed ASL.

In addition to accounts that are motivated by predictive operations in language processing and their consequences for perceptuo-motor activations, as above, there is intriguing imaging evidence for language processing proper being executed in early visual areas, in

congenitally blind adults (Bedny et al. 2011; Bedny & Saxe, this volume). They show that phonological, semantic, and syntactic information is differentially processed in visual cortex in these participants, suggesting that these areas can be recruited in addition to classical areas of the language network. Whatever the properties of the local visual neural circuits are, they can operate over linguistic representations in a way that distinguishes language processing from other tasks, suggesting that the computations that comprise language processing are sufficiently generic to interface with those types of circuits.

What are we to conclude from such examples? Both classical language areas (Broca's region) and putatively domain general areas (early visual cortex) are implicated in different forms of language tasks. Critically, it is not just the visual form of the materials that drive the low-level activation. In the cases mentioned here, early visual cortex is (re) organized to reflect sensitivity to high-order properties of high-level category information in one case and human-form-based analysis in the other. In neither case do feedforward analyses account for the patterns; rather, predictive aspects of processing, based on knowledge of language, appear to condition the neural response properties. Similarly, (the sub-regions of) Broca's region are not just driven by high-level language constructs; relatively low-level discriminations implicate Broca's region as well. Presumably it follows from this discussion of maps of regions that we must decompose linguistic computation into constituent elements that are much finer (and perhaps of a different kind) than we are used to doing. For instance, arguing that a brain region forms the basis for 'syntax' or 'phonology' fails to capture the computational structure that actually comprises something as richly structured as syntax (which is obviously not monolithic). Surely we will not want to argue that, say, visual cortex 'hosts' syntax. But we are then required (as linguists, psychologists, cognitive neuroscientists) to provide an analysis that captures which component of syntactic computation in fact is compatible with what is executed in visual cortex. In the cases sketched out above, aspects of predictive coding that lead to sensorimotor predictions are the likely suspects.

Maps of streams

In vision, the concept of processing streams has been immensely influential in shaping research since the 1980s. The existence of concurrent *what*, *where*, and *how* streams showed how certain large-scale computational challenges (localization, identification, sensorimotor transformation/action-perception linking) could be supported in parallel by streams consisting of hierarchically organized cortical regions. This architecture provided a functional anatomic blueprint for supporting parallel processes as well as the hierarchical elaboration of visual representations within a stream (Mishkin et al. 1983; Goodale & Milner 1992).

This key idea from systems neuroscience was adopted and adapted for speech and language processing in the last ten years (Hickok & Poeppel 2004, 2007; Friederici 2012; Rauschecker & Scott 2009). Figure 2 illustrates one such model, constructed to account for a range of phenomena in speech processing. A ventral stream can be considered responsible for supporting the transformation from auditory input to lexical representations. A dorsal stream is primarily critical for sensorimotor transformations. One noteworthy feature, and a clear departure from previous models, is that current data suggest much more bilateral support for speech and language processing than was previously suggested. The ventral stream structures, in particular, are robustly bilateral.

As is true for vision, identifying objects is one fundamental computational challenge in speech perception. The other computational role (in speech), in the sense of Marr, is to provide the infrastructure for production. That is to say, the words we know must be able to make contact with both articulatory representations and auditory representations. The dorsal

stream regions appear to host the operations that permit the seamless translation from auditory type (e.g. spectrogram) to articulatory type (e.g. motor code) representations.

Note how both of these computational challenges link to some foundational questions of speech acquisition. If the characterization of such processing streams is on the right track, developmental data will illuminate the role of these structures. For example, consider that the identification of speech sounds and early lexical processing far precedes their articulation. In fact, performance in relatively complex language tasks – such as mapping novel word forms to visual referents, exploiting some remarkably high level intonational cues – can be documented by 6 months of age (Shukla et al. 2011). Does that mean that the ventral stream is mature prior to the dorsal stream? More provocatively, does this mean that articulatory information is not causally necessary for the establishment of phonological and lexical representations? Is dorsal stream information merely modulatory? That speech motor information plays an important modulatory role in perception is not under debate; however, data on the development of these processing streams will sharpen the questions about what information is necessary and sufficient for establishing the speech system, in the learner.

Maps of hemispheres - lateralization of function

Zooming even further out (than regions and streams), arguably the coarsest level of brain mapping is at the level of cerebral lateralization. However, coarse or not, hemispheric asymmetry has been an issue at the very foundation of cognitive neuroscience, and particularly in the context of language processing. Since the very first insights by Broca, Wernicke, Lichtheim, and others, the strong lateralization of ‘language’ processing to the dominant hemisphere has been one of the hallmarks of brain organization.

If the analysis proceeds at the level of ‘language,’ as if it were an unanalyzed whole, an equivalently coarse biological analysis is commensurate – and well deserved. As soon as one begins to decompose language into its constituent domains and operations, the question of cerebral asymmetry becomes more complicated and nuanced. For example, there is now emerging consensus that speech perception – roughly the mapping from acoustic waveform to lexical representation – is a bilaterally mediated subroutine of language comprehension (cf. Figure 2). There is also a growing body of evidence to suggest that some aspects of lexical level processing (especially lexical semantics) are executed in both hemispheres (Federmeier et al. 2008). In contrast, syntactic processing and production appear to be rather lateralized to the dominant hemisphere. To caricature the pattern, the linguistic operations that appear to be robustly associated with left hemisphere mechanisms are “COM-COM-PRE-PRO”: combinatorics, composition, prediction, and production. Notice that these terms refer to operations or computations, not to representational primitives. To date there exists no compelling evidence that, say, distinctive features, or morphemes, or roots, or phrasal types are selectively lateralized. Stored linguistic information may be encoded in various cortical regions in the left and right hemispheres, but the computations that operate over the putative representations appear to reflect lateralized specializations.

One area of investigation in which there has been vigorous debate concerns the relative contributions of left and right hemispheres to speech perception proper. Given that speech perception appears to be bilaterally mediated, what are the two sides doing? Are they fully redundant? Completely different? Are there relative specializations?

To get a flavor for some of the recent hypotheses, it is helpful to recall what the perceptual system must ultimately achieve, with speech or other auditory inputs. On the one hand, the system must extract sequential items at a time resolution that permits identifying the correct order of parts. Words, in particular, are comprised of a sequence of relatively short segments, and correctly identifying their order is a prerequisite for recognition (the *pest*

versus *pets* issue). The high time resolution necessary to succeed demands neural circuits optimized to integrate information over short time scales (on the order of tens of milliseconds). On the other hand, listeners must be able to analyze and use small frequency deviations, for example in the context of suprasegmental prosody (the *lunch!* versus *lunch?* issue). Such changes demand a high frequency resolution and are often associated with longer time constants. Again, neural circuitry must be responsive to slower time constants and high spectral sensitivity. The output of both types of perceptual analyses must be in a form that can interface with the language system, let us say for the sake of argument, lexical representations.

It has been suggested, in the context of two slightly different but closely related theories (Zatorre et al. 2002; Poeppel 2003), that right hemisphere auditory mechanisms show a higher sensitivity to spectral change and are optimized for longer time constants, whereas left lateralized mechanisms may reflect better temporal resolving power. The spectral versus temporal theory by Zatorre and colleagues and the asymmetric sampling in time (AST) theory by Poeppel appeal to slightly different attributes of cortical neurons, but the models by and large agree that some type of split must be performed (Zatorre & Gandour 2008). It is fair to say that there is emerging consensus about right hemisphere properties. Whether not the left hemisphere is optimized for rapid temporal integration or short temporal sampling remains quite controversial (e.g. McGettigan & Scott 2012). To test whether such asymmetries were *intrinsic* properties of the auditory system (and not just associated with processing overt stimuli), Giraud and colleagues (2007) recorded concurrent EEG and fMRI data while subjects were resting (no stimulus input). They filtered the EEG data into bands commensurate with the time scales of interest (~ theta, ~ gamma) and used that electrophysiological pattern (convolved with the hemodynamic response function) to query the fMRI data. In two separate studies, the same lateral asymmetry was observed: the higher modulation rate gamma band responses correlated with left-hemisphere auditory structures, the lower rate theta responses with the right auditory cortex. These data demonstrate that *there exists some intrinsic asymmetry such that auditory neuronal ensembles 'hear' the world at different time scales simultaneously*. This (implementational) asymmetry provides the basis for the (algorithmic) multi-time resolution approach that solves the (computational) integration-resolution tension for auditory signals.

Maps of the whole network

Recently, the neurosciences have witnessed substantial enthusiasm for trying to map across local as well as distant areas, or even the entire brain, using increasingly sophisticated network analyses (e.g. Sporns 2010), and building on the new techniques and their improved ability to reveal connections at both fine and coarse scales. The appetite for this type of work culminates in connectomics (Seung 2011), the presupposition of which is that understanding all the connections will lead to understanding of how the mind works.

Attempting to map the connections of the brain is, surely, a laudable goal. What lies behind this surge in research effort? There are at least two strands of work that fuel the desire to map everything - and connect everything. One motivating factor presumably derives from the remarkable insights gleaned from the cutting edge techniques in neuroscience. *Two-photon microscopy* (Homma et al., 2009) and *optogenetics* (Yizhar et al., 2011) are generating stunning new data on local circuitry and connectivity. These are, it should go without saying, not techniques easily applied to undergraduates or babies... They are, however, unveiling in animal preparations obscenely high-resolution views on the structure and function of small local regions of central nervous system. If it is possible to map local circuits with such clarity—and, indeed, such aesthetic beauty—then it stands to reason that would be exciting to scale up such approaches to whole brains. A second motivating factor appears to lie in the belief that acquiring enormous amounts of data and mining them

statistically will yield new types of insights. There is an unending thirst for more and more data (the same cannot be said for more and more theory) (For an engaging perspective on the problem of data collecting versus theory constructing, in physics and in linguistics, see Weinberg 2009 and Chomsky 2000, respectively). But beyond being numerically awesome, will a comprehensive map, the connectome of the human brain, form the basis for the explanation of the human mind? Will the wiring diagram suffice? It seems unlikely. For example, the worm *C. elegans* has a nervous system that is completely known and described, yet our understanding of the function is embarrassingly inadequate (Bargmann 1993).

It is therefore possible that that even something as sophisticated and data-driven as the connectome will be at best an intermediate step - unless we also make serious, theory-driven progress on the cognome (the specification of the representational and computational parts list of human cognition) and function. Genomics has been a triumph of computational biology in terms of capturing something truly numerically awesome, but its real utility won't be revealed until we understand the *coding* at a much more comprehensive level (including, say, the proteome).

In conclusion, can making maps or defining topography ever meet the criterion of explanatory adequacy? Spatial maps are not explanations, so there is an important ingredient missing, a detailed functional analysis. We need to begin to address the difficult link between maps and computation. This presumably requires computational theories of a certain granularity; that is to say, what is localized are computations for which a certain circuit is specialized (perhaps 'canonical computations,' as recently discussed by Carandini, 2012). The maps we make will therefore be maps of computations.

The maps problem and developmental data

A variety of developmental questions intersect with the map perspective outlined above (admittedly, from the naïve point of view of a non-developmentalalist). I describe just three in brief below, focusing on data that speak to three of the four map types. However, it is worth emphasizing that data collection in a developmental context will also not suffice without theory building. In particular, computational theories will provide the most useful framework for addressing the development of the cognitive and neural foundations of perception and cognition (Aslin & Fiser 2005).

Figure 1 depicts the organization of the adult Broca's region subsequent to normal language acquisition. An obvious question is whether the infant brain similarly structured. More specifically, are such cytoarchitectural subdivisions consequences of the acquisition of speech and language (and other aspects of cognition) – or causes? Are different subregions more or less susceptible to modification by experience? Can we use data from developmental cognitive neuroscience to begin to unpack the potential contribution of distinct regions? Given the enormous appetite for data from the infant brain – including for speech and language processing (e.g. Dehaene-Lambertz et al. 2006; Telkemeyer et al. 2009; Perani et al. 2011) – one viable research strategy is to use knowledge of development to understand features of the brain (rather than the other way around).

Turning from Broca's region to the dual pathways, there is now a growing body of fiber tracking evidence (largely from diffusion tensor imaging) that suggests that in two-day old infants, many ventral pathway structures are in place in an adult-like configuration, while the dorsal stream is anatomically underspecified (Perani et al. 2011, Friederici 2012). Figure 4 illustrates some relevant data. This pattern is certainly consistent with an interpretation that dorsal stream regions have a critical role for production, or the linking of representations from temporal to frontal regions by sensorimotor transformation. Note that there is more than one dorsal stream projection depicted in Figure 2. It has been argued that

the projection targeting Broca's region is not in place, whereas the more dorsal structures are available earlier in ontogeny. If this finding is correct, it reinforces the issue raised in the section above on Broca's region: does the fine-grained spatial structure visible there arise in development with the 'arrival' of the fiber tracts from temporal lobe origins – and can we document concomitant performance changes – or is it in place and is merely modified as the fiber tracts mature? Many more normative data sets on the developing brain will be necessary to sort out such issues.

Finally, consider hemispheric asymmetries. There now exists a growing body of evidence that right hemisphere mechanisms are developmentally slightly advanced. This developmental asymmetry appears to be the case both for the auditory and visual systems. Given that, by hypothesis, the local circuitry in the right hemisphere is different with respect to its processing preferences (not necessarily different in kind, just in allocation of amount of neural ensembles with certain properties, say slow temporal integration), this implementational difference provides a way to address a few perceptual phenomena that require explanation. For example, in the auditory domain, it would be plausible for infants to show enhanced sensitivity to syllabically mediated information as well as other, suprasegmental prosodic phenomena, even in utero. The 'universal' listener may start out with a syllable-based segmentation strategy. In general, larger units, preferentially processed by right hemisphere mechanisms, would be favored in early development. Similarly, configural perceptual information should be preferred to locally detailed, high-resolution cues.

Given some of the properties that have been discussed concerning hemispheric asymmetry, the dorsal and ventral streams, and the anatomic specification of Broca's region, one can advance a (simplified) developmental hypothesis arising from these cognitive neuroscience descriptions. (i) The right-hemisphere mechanisms, especially in the superior temporal lobe portions of the ventral stream, are available slightly (weeks to few months) earlier in ontogenesis, perhaps even in utero. Given the hypothesized specialization of right temporal lobe mechanisms for longer integration time constants, the perceptual information attended to preferentially by the learner occurs on these longer timescales. Units of approximately syllabic size/duration – which corresponds closely to the envelope of spoken language – have epistemological priority at this stage; perceptual analysis at this temporal granularity is especially effective at helping the learner segment the input into units that correspond well to certain linguistic grouping principles (stress, consonant-vowel alternation, possibly even rhythm class) (cf. Bertoncini & Mehler 1981).

(ii) Convergence by the learner on the phonetic repertoire of her language over the first year of life (see, e.g. Kuhl 2004 for review) coincides with the increasing specification of the structures underlying segmental phonology. Since left hemisphere mechanisms have been argued to reflect some degree of specialization for short time scale processing, and dorsal stream structures are implicated in processing segmental phonology, this accords well with the requirement to analyze segmental and subsegmental phonetic information. On balance, it is plausible to assume that the ventral stream structures are in place by approximately one year. The learner is now optimally positioned to analyze input both at syllabic and segmental timescales, concurrently. This combination constitutes critical ingredients for successful decoding of words – the second big milestone the learner has to achieve (i.e. the incremental construction of the mental lexicon). That is not to say that younger infants cannot discriminate segmental information or fail to map auditory word forms to visual referents (see, e.g., Shukla et al. 2011). These tasks, though, could be achieved by the learner based on representations that are underspecified with respect to the linguistic inventory that is ultimately specified for their native language; that is to say, more acoustically driven,

episodic representations could form the basis for executing such tasks in the first few months of development.

(iii) The structures of the dorsal stream—or at least one of the dorsal streams—are argued to develop later than the ventral stream. The dorsal stream projection from the temporal lobe to pre-motor areas has been suggested to be present at birth. In contrast, the dorsal projections to the inferior frontal gyrus, including Broca's region, show a delayed developmental trajectory (Perani et al. 2011, Friederici 2012). The structures of the dorsal stream play a critical role in sensorimotor transformation for speech. The seamless interaction of the perception and production systems is arguably coordinated by dorsal stream structures, and it is therefore plausible that there is a tight association between the neural development of the dorsal stream and the developments of speech production. Uncontroversially, perception precedes production, and it is reasonable to hypothesize that this delay in part reflects the constraints imposed by a developing sensorimotor dorsal stream.

(iv) If inferior frontal gyrus regions, and in particular Broca's region, receive a changing innervation pattern over development because the dorsal stream projecting there is not fully developed, it is reasonable to conjecture that the functional specification of some of these regions changes (perhaps sharpens) over time. It is not possible to conclude whether or not the many cytoarchitectural subregions of the IFG (Figure 1) are available at birth, or whether this region becomes increasingly specified over development. One can, however, conclude that the computational infrastructure is still changing over time, as the inputs to this rich and complex structure change quite dramatically. A thoughtful analysis of the development of speech and its cerebral lateralization is provided by Minagawa-Kawai et al. (2011). They discuss the interaction between signal-driven lateralization versus lateralization as a function of learning biases and derive a model of the acquisition of speech that covers many of the phenomena raised here. It is worth noting that several of the subregions of the IFG play critical functional role in the processing of complex sentential representations, as well as in other domains of language and non-language processing. (Similarly, the thorny issues of semantics remain untouched in this conjectured scenario.) The sequence outlined here is narrowly focused on the possible development of speech processing.

The future of cognitive neuroscience: making mappings

We now turn from making maps, which lies at the very foundation of cognitive neuroscience, to the problem of *mappings*. This latter problem refers to a very different way of thinking about the challenges that the field faces. What is at stake is whether we are able to identify the possible formal *relations* (mappings) between the putative primitives of cognitive science, and in particular language research, and the putative primitives of neurobiology. Whereas the maps problem raises a number of practical questions – for example, what might be the right techniques to generate topographic maps at the right level of analysis – the mapping problem raises principled questions about the nature of the relation between psychology and neuroscience. This issue has been addressed in some detail by Poeppel & Embick (2005) for language and is, of course, raised by many investigators who grapple with the challenges at the interface of mind and matter (Smolensky & Legendre 2006).

Granularity mismatch and the ontological incommensurability: misaligned parts lists

One way to develop an intuition for the problem is to think about it in the context of specifying the 'parts lists' of the mind and brain. More formally speaking, this refers to the goal of specifying the *set of primitives*, or the ontological structure, of the two domains in question. As cognitive scientists, we aim to determine an exhaustive list of the components of the human mind—call it *the human cognome*. We use the inventory of components to

explain fundamental features of perception and cognition; this is done by appealing to the hypothesized primitives and the rules that govern their interaction to yield the phenomena that we are investigating. This is no different in kind from physics, where the field aims, for example, to specify the elementary particles as well as the forces that condition their interaction—and these two features of a theory provided an explanation of the phenomena. Similarly, as neuroscientists we seek to determine the infrastructure of the nervous system, from sub-cellular components to large-scale systems.

Let us make the problem a bit more concrete, building on the notion of the parts list. Suppose we write to a large (even exhaustive) list of speech and language researchers, asking what they consider to be the primitive elements of their domain without which they could not account for the elementary phenomena of their fields. In other words, what are the set of primitives that are absolutely required to provide an explanatory account of the canonical phenomena in speech perception, language comprehension, and production? An answer to such a query will yield a *parts list of representational elements* as well as a *list of elementary functions, or computations*. Figure 5 (left column) illustrates a hypothetical answer. An enumeration of primitives might include concepts such as distinctive feature, syllable, morpheme, small clause, variable, etc. The elementary functions might include concatenation, recursion, variable binding, etc. Note that such a list will be theory-dependent. A particular theory might dispute the fundamental nature of distinctive features or the relevance of recursion. However, specific commitments of a given theory are beside the point; *any* theory (of linguistics, psycholinguistics, neurolinguistics) must be able to provide an inventory of its constituent representations and computations. Obviously, it is of particular interest to develop a consensus position on what the ‘linguistic cognome’ must necessarily contain to begin to get a grip on what is the ontological infrastructure of language research.

Now, this very same exercise can be performed with neuroscientists. The right column in Figure 5 illustrates the type of responses one might expect. The goal there, too, is to identify an emerging consensus about the neurobiological units that form the basis of central nervous system structure and function. For the sake of argument, Figure 5 now provides a *parts lists for the language sciences* and a *parts list for the neurosciences*. The evidence for the primitives is extensive. That is to say, there is a rich body of empirically-based as well as theoretical argumentation to support every given member of these two “alphabets.”

The mapping problem is now quite obvious: what is the relation between the ontological structure of the language sciences and the neurosciences? More colloquially, what are the lines/arrows that one can draw from one side to the other in Figure 5 that can capture the systematic mapping from one set of elementary units to another? The answer, too, is obvious: *there is absolutely no mapping to date that we understand in even the most vague sense*. There are no equivalence relations, no isomorphisms, no easy mappings from the theoretical infrastructure of the cognitive sciences to that of neurobiology. On the assumption that one is adopting a non-dualist research strategy, it is entirely unclear how to link the two approaches (beyond stating some relatively gross correlations). In what follows, a brief diagnosis for this state of affairs is provided. Then a suggestion is made, in the spirit of Marr (1982) and Gallistel & King (2009) about how to move forward, adopting an optimistic stance about computational cognitive science and computational neuroscience. The central question will be: what are plausible linking hypotheses between the two domains, what form can they (or must they) take? Can cognitive neuroscience of speech and language processing move beyond a purely correlational state of affairs?

There are (at least) two reasons why the formulation of linking hypotheses has been slow – beyond the fact that it is rarely if ever acknowledged that there is a problem to begin with.

One reason that it has proven difficult derives from the level of analysis that is typical of the different fields. Linguistic research comes at the questions of interest with a toolbox that permits a very high 'resolution.' An area of research as conceptually straightforward as *word recognition* approaches the issue by incorporating subtle features of phonology, the statistical adjacency relations between sounds, the role of morphological structure, the role syntactic categories may play for lexical access, nuanced theories about the semantics of words, theories exploiting the predictive nature of language processing, etc. to understand aspects of facilitated or impeded lexical processing, and so on. That is to say, even a process as *allegedly* straightforward as recognizing a spoken word interfaces in a variety of ways with phonology, morphology, syntax, lexical semantics, compositional semantics, and aspects of discourse. In contrast, neurobiological research on language attempts to address much broader questions. Typical issues that are investigated in imaging studies, for example, include whether syntax is localized or where the mental lexicon is localized. This is, to be sure, not a problem that cannot be overcome. It is a practical challenge and simply reflects what has been called the *granularity mismatch problem* (Poeppel & Embick 2005), to highlight the fact that the granularity of analysis is too different to permit links that are more than correlational. The techniques available to the brain sciences investigate the issues in a manner commensurate with what can be measured and analyzed; the techniques available to the cognitive sciences attempt to capture generalizations about speech and language processing at the highest possible 'conceptual resolution.' The practical consequence of such a granularity mismatch is that issues at very different levels of representation end up being addressed. This is not intrinsically bad, but it is limiting in that it fails to move the field forward in developing mechanistic linking hypotheses between brain structure and function and the organization of this particular cognitive system. In recent years, this has been changing in a positive direction in the context of the ever-closer relations between theoreticians, cognitive scientists, and neurobiologists applying state-of-the-art techniques.

In contrast to the granularity mismatch problem, which can be overcome by forging closer links between language researchers and brain researchers to better coordinate the nature of the questions can be profitably addressed, there is another consideration that is a bit more insidious, and for which a solution is more elusive. We seek to account for the phenomena of cognition by appealing to the neurobiological infrastructure that lies at its basis. However, as is illustrated in Figure 5 and intimated above, there is no comprehensible way in which we can *link* the two sides beyond correlation - and this may reflect more than just a mismatch in the granularity of analysis. It may reflect a principled incommensurability. The *ontological incommensurability problem* (Poeppel & Embick 2005) suggests that the primitive elements of the two domains cannot be mapped onto each other at all given the current formulation of the parts lists. On such a view, the concepts that form the basis of cognition and neurobiology are impossible to align in principle. If that is the case, then there is no opportunity for identifying plausible isomorphisms and certainly no opportunity for reduction (if that were a goal). Given—plausibly—that there is no hope to reduce perception and cognition to neurobiology at present (and in the most primitive sense of reductionism), an alternative goal is to aim for unification (Chomsky 2000) or consilience (1999). The larger question thus becomes: what steps need to be taken in terms of how we talk about a given domain in order for it to be unified with a different domain? A historical example that illustrates such a case comes from the relation between chemistry and physics at the beginning of the 20th century. Chemistry was, contrary to intuition, not reduced to the principles of physics. Rather, the conceptual structure (ontology) of physics had to change to adjust to the new insights coming from the (putatively higher order) area of chemistry. Provocatively, a similar situation might present itself in the relationship between the brain and cognitive sciences (Chomsky 2000, Gallistel & King 2009). Perhaps what are considered elementary functional units in the neurosciences will change by considering more closely what the demands of perception and cognition are. Conceivably, a

representational primitive such as ‘neuron’ may not end up playing as foundational role as some other functional units yet to be determined, even if it is presumably constructed out of the existing parts (say, for example, a canonical cortical microcircuit, Douglas and Martin 1991).

To summarize, the *granularity mismatch problem* is a practical problem, a pragmatic approach to which will play a productive role in dealing with the maps problems addressed in the first section. The *ontological incommensurability problem* reflects a principled problem, and to make progress on this challenge, practitioners may need to consider how to reformulate the nature of interdisciplinary research in this area - and question the commitments we make to the primitives of representation and computation.

A tentative strategy for progress

One approach to explore can be called the ‘radical decomposition’ strategy. Adopting such a strategy means questioning whether the primitives we hold to be foundational are possibly decomposable into finer-grained elements (that may themselves not seem like natural units at first glance). For example, from the perspective of the cognitive sciences, consider a concept such as the ‘phoneme.’ Many decades of research show that the phoneme is a bundle of more elementary units (features) and that generalizations about knowledge of phonology are not made at the level of phonemes but at the level of features. Similarly, consider the tension between a theoretically poorly motivated concept such as ‘word’ compared to a theoretically richly supported concept such as ‘morpheme.’ Whereas morphemes play an essential role in accounting for a range of properties of lexical structure and lexical access, an informal notion such as ‘word’ has been largely unsuccessful. By analyzing traditional concepts at a higher resolution, the field has successfully identified smaller (or different) units that appear to do a better job at capturing the phenomena under investigation.

What kind of smaller – or different - units might be plausible? One strategy is to identify representations and operations that can be linked to the types of operations which simple electrical circuits can execute. One might begin, for example, with theoretically well-motivated units of representation or processing deriving from cognitive science research (here, say linguistics); then one attempts to decompose these into elementary constituent operations that are formally generic (something like ‘concatenation,’ for example). Now, in a role reversal, linguists should challenge neurobiologists to define and characterize the neural circuitry that can underpin something as elementary as concatenation. Instead of seeking validation of the hypothesized units from cognition by putative reduction to neuroscience, assume that the cognitive evidence is *as foundational* and motivate research to seek a neural circuit can execute the type of operation known to be necessary for a variety of cognitive tasks that are beyond debate. Why might such a (more muscular) stance of cognition be helpful? The field has taken for granted that the elementary units of neurobiology constitute the fact of the matter, the ground truth, from a functional point of view. Neurons, dendrites, dendritic spines, cortical columns, long-term potentiation... The extremely impressive parts list of the neurosciences is not under debate. What one might question, though, is whether the arrangement of units as they are currently discussed reflects the functionally appropriate arrangements to account for the phenomena we know to require explanation. The interaction between the evidence coming from cognitive science theory and experimentation, on the one hand, and the practice and interpretation of neurobiology, on the other, needs to be significantly more synergistic to go beyond correlation and develop genuinely explanatory models.

At this point it is worth considering once more the research strategy outlined by Marr and others. The computational level of analysis, provided by linguistics, psychology, and aspects

of computer science, must be linked to the implementational level of analysis, neurobiology. What was suggested by Marr is that an intermediate (algorithmic) level of description specifies the representations and computations that are executed by the implementational circuitry to form the basis of the computational level of characterization. If this strategy is on the right track, current research should in part focus on the operations and algorithms that underpin language processing. The commitment to an algorithm or computation in this domain commits one to representations of one form or another with increasing specificity and also provides clear constraints for what the neural circuitry must accomplish. The kinds of operations that might provide the basis for investigation include concatenation, segmentation, combination, labeling, and other elementary (and generic) operations that could be implemented quite straightforwardly in neural circuits.

A developmental wedge

On the view sketched here, a prerequisite to develop principled linking hypotheses between (linguistic) cognition and neurobiology is the definition of the parts list, or the cognome. Developmental research has a central role to play in this research endeavor. No other area of psychological investigation deals more explicitly with the question of what the elementary constituents of the mind are. An essential perspective on the question of the ontological primitives comes from the developmental literature, and specifically the question “what is the parts lists at the start”? It would be of tremendous value to the enterprise of mapping from cognition to neurobiology if developmental approaches could integrate questions about the computational constituents. For example, instead of framing a developmental linguistic question in terms of the domain (e.g. speech, phonology) or the tasks employed (e.g. is phonetic categorization evident in the neonate?), if such questions could be characterized in terms of their elementary constituents it would license a claim about the elementary computational infrastructure. Much as in adult cognitive neuroscience research, describing the nature of the perceptual and cognitive tasks not as tasks but as a set of primitive operations and representations will allow the formulation of linking hypotheses that translate more seamlessly to computational neuroscience.

There are, to be sure, good examples of the direction that can be taken. A lot of excellent work in vision research demonstrates how tightly one can link theoretically motivated and computationally explicit theories of visual perception to detailed neural models that link these domains. My own favorite example comes from the auditory research area of sound localization. The avian (e.g. barn owl) versus mammalian (e.g. hamster) brain solve the same sound localization problem (computational level of analysis) using different algorithms (delay line model versus phase) using specialized circuits that implement only one kind of computation and not another – that is to say, the circuits explain the computation (Grothe 2003). The overall goal, presumably, remains to provide *explanations* of how the properties of the brain account for perception and cognition - not just *descriptions*. The state of the art is impressive, and perhaps even descriptively adequate, but our accounts of the relation between brain and cognition remain in large part correlational. Cognitive neuroscience ought to be more ambitious than that, striving for explanations. A big issue in this context is whether we are investigating cognition at the right granularity, if our goal is to develop linking hypotheses to neurobiology. If we stick to cognition per se, and we achieve adequate explanations, then perhaps the analyses are on the right track. However, a more radical decomposition approach might be necessary to develop successful linking hypotheses.

Many of the issues that lie at the center of current research are extremely relevant, and only a slight shift in perspective can provide a different view on how they help address the linking challenge. For example, making maps, or localization of function, continues to be extremely influential; but localization of function might better be spelled out as localization

of local computation - where computation is now construed to be the mid-level, generic set of operations that underlie cognitive representation of computation. Similarly, the relentless enthusiasm for prediction in perception and cognition, now largely subsumed under Bayesian approaches, exemplifies the types of elementary computations that might hold in different parts of perception and cognition, linguistic or otherwise.

References

- Almeida D, Poeppel D, Corina D. The brain tuning to the body: Cortical plasticity and the perceptual encoding of gestures. under review.
- Amunts K, Lenzen M, Friederici AD, Schleicher A, Morosan P, Palomero-Gallagher N, Zilles K. Broca's region: Novel organizational principles and multiple receptor mapping. *Plos Biology*. 2010; 8(9):10.1371/journal.pbio.1000489
- Anwander A, Tittgemeyer M, von Cramon DY, Friederici AD, Knösche TR. Connectivity-Based parcellation of broca's area. *Cerebral Cortex* (New York, N.Y : 1991). 2007; 17(4):816–25.10.1093/cercor/bhk034
- Aslin RN, Fiser J. Methodological challenges for understanding cognitive development in infants. *Trends in Cognitive Sciences*. 2005; 9(3):92–98. [PubMed: 15737816]
- Bargmann CI. Genetic and cellular analysis of behavior in *C. Elegans*. *Annual Review of Neuroscience*. 1993; 16(1):47–71.
- Bedny M, Pascual-Leone A, Dodell-Feder D, Fedorenko E, Saxe R. Language processing in the occipital cortices of congenitally blind adults. *Proceedings of the National Academy of Sciences*. 2011; 108(11):4429–34.
- Bedny M, Saxe R. *Cognitive Neuropsychology*. in press.
- Bertoncini J, Mehler J. Syllables as units in infant speech perception. *Infant Behavior and Development*. 1981; 4(0):247–260.10.1016/S0163-6383(81)80027-6
- Broca P. Perte de la parole, ramollissement chronique et destruction partielle du lobe antérieur gauche du cerveau. *Bull Soc Anthropol*. 1861; 2:235–238.
- Carandini M. From circuits to behavior: A bridge too far? *Nature Neuroscience*. 2012; 15(4):507–9.10.1038/nn.3043
- Carter, R.; Frith, CD. *Mapping the mind* (illustrated, reprint ed.). Berkeley: University of California Press; 2000.
- Chomsky, N. *New horizons in the study of language and mind* (illustrated, reprint ed.). Cambridge, UK; New York: Cambridge University Press; 2000.
- Dehaene-Lambertz G, Hertz-Pannier L, Dubois J. Nature and nurture in language acquisition: Anatomical and functional brain-imaging studies in infants. *Trends in Neurosciences*. 2006; 29(7): 367–73.10.1016/j.tins.2006.05.011 [PubMed: 16815562]
- Dikker S, Pykkänen L. Before the N400: Effects of lexical-semantic violations in visual cortex. *Brain and Language*. 2011; 118(1-2):23–8.10.1016/j.bandl.2011.02.006 [PubMed: 21458057]
- Dikker S, Rabagliati H, Pykkänen L. Sensitivity to syntax in visual cortex. *Cognition*. 2009; 110(3): 293–321.10.1016/j.cognition.2008.09.008 [PubMed: 19121826]
- Dikker S, Rabagliati H, Farmer TA, Pykkänen L. Early occipital sensitivity to syntactic category is based on form typicality. *Psychological Science : A Journal of the American Psychological Society / APS*. 2010; 21(5):629–34.10.1177/0956797610367751
- Douglas RJ, Martin KA. A functional microcircuit for cat visual cortex. *The Journal of Physiology*. 1991; 440:735–69. [PubMed: 1666655]
- Embick D, Poeppel D. Mapping syntax using imaging: Problems and prospects for the study of neurolinguistic computation. *Encyclopedia of language and linguistics*. 2006
- Federmeier KD, Wlotko EW, Meyer AM. What's "right" in language comprehension: Erps reveal right hemisphere language capabilities. *Language and Linguistics Compass*. 2008; 2(1):1–17.10.1111/j.1749-818X.2007.00042.x [PubMed: 19777128]

- Friederici AD. The cortical language circuit: From auditory perception to sentence comprehension. *Trends in Cognitive Sciences*. 2012a; 16(5):262–8.10.1016/j.tics.2012.04.001 [PubMed: 22516238]
- Friederici AD. Language development and the ontogeny of the dorsal pathway. *Frontiers in Evolutionary Neuroscience*. 2012b; 4:3.10.3389/fnevo.2012.00003 [PubMed: 22347185]
- Gallistel, CR.; King, AP. Memory and the computational brain why cognitive science will transform neuroscience. Hoboken: John Wiley & Sons; 2011. Retrieved from <http://public.eblib.com/EBLPublic/PublicView.do?ptiID=819438>
- Geschwind N. The organization of language and the brain. *Science*. 1970; 170:940(3961)
- Giraud AL, Kleinschmidt A, Poeppel D, Lund TE, Frackowiak RS, Laufs H. Endogenous cortical rhythms determine cerebral specialization for speech perception and production. *Neuron*. 2007; 56(6):1127–34.10.1016/j.neuron.2007.09.038 [PubMed: 18093532]
- Goodale MA, Milner AD. Separate visual pathways for perception and action. *Trends in Neurosciences*. 1992; 15(1):20–25. [PubMed: 1374953]
- Grodzinsky, Y.; Amunts, K. Broca's region (illustrated ed.). Oxford; New York: Oxford University Press; 2006.
- Grothe B. New roles for synaptic inhibition in sound localization. *Nature Reviews Neuroscience*. 2003; 4(7):540–50.10.1038/nrn1136
- Hickok G, Poeppel D. The cortical organization of speech processing. *Nature Reviews Neuroscience*. 2007; 8(5):393–402.10.1038/nrn2113
- Hillis AE, Caramazza A. Representation of grammatical categories of words in the brain. *Journal of Cognitive Neuroscience*. 2007;10.1162/jocn.1995.7.3.396
- Homma R, Baker BJ, Jin L, Garaschuk O, Konnerth A, Cohen LB, Zecevic D. Wide-Field and two-photon imaging of brain activity with voltage-and calcium-sensitive dyes. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2009; 364(1529):2453–2467.
- Hubel, D. Eye, brain, and vision (2, reprint, illustrated ed.). New York: Scientific American Library; 1995.
- Kuhl PK. Early language acquisition: Cracking the speech code. *Nature Reviews Neuroscience*. 2004; 5(11):831–843.
- Marr, D. Vision: A computational approach. San Francisco: Freeman & Co; 1982.
- Mausfeld R. On some unwarranted tacit assumptions in cognitive neuroscience. *Frontiers in Psychology*. 2012; 3:67.10.3389/fpsyg.2012.00067 [PubMed: 22435062]
- McGettigan C, Scott SK. Cortical asymmetries in speech perception: What's wrong, what's right and what's left? *Trends in Cognitive Sciences*. 2012; 16(5):269–76.10.1016/j.tics.2012.04.006 [PubMed: 22521208]
- Minagawa-Kawai Y, Cristia A, Dupoux E. Cerebral lateralization and early speech acquisition: A developmental scenario. *Developmental Cognitive Neuroscience*. 2011; 1(3):217–32.10.1016/j.dcn.2011.03.005 [PubMed: 22436509]
- Mishkin M, Ungerleider LG, Macko KA. Object vision and spatial vision: Two cortical pathways. *Trends in Neurosciences*. 1983; 6:414–417.
- Obleser J, Elbert T, Lahiri A, Eulitz C. Cortical representation of vowels reflects acoustic dissimilarity determined by formant frequencies. *Cognitive Brain Research*. 2003; 15(3):207–213.10.1016/S0163-6383(81)80027-6 [PubMed: 12527095]
- Perani D, Saccuman MC, Scifo P, Anwander A, Awander A, Spada D, Friederici AD, et al. Neural language networks at birth. *Proceedings of the National Academy of Sciences of the United States of America*. 2011; 108(38):16056–61.10.1073/pnas.1102991108 [PubMed: 21896765]
- Pinker, S. How the mind works (illustrated, reissue ed.). New York: Norton; 1997.
- Poeppel D. The analysis of speech in different temporal integration windows: Cerebral lateralization as 'asymmetric sampling in time'. *Speech Communication*. 2003; 41(1):245–255.
- Poeppel, D. Bildwelten des Wissens: Vol 6.1. Bildwelten des wissens. Berlin: Akademie Verlag; 2008. The cartographic imperative: Confusing localization and explanation in human brain mapping; p. 13

- Poeppel D, Embick D. Defining the relation between linguistics and neuroscience. *Twenty-First Century Psycholinguistics: Four Cornerstones*. 2005:103–118.
- Poeppel D, Hickok G. Towards a new functional anatomy of language. *Cognition*. 2004; 92(1):1–12. [PubMed: 15037124]
- Posner, MI.; Raichle, ME. *Images of mind* (2, illustrated, reprint ed.). New York: Scientific American Library; 1997.
- Pulvermüller F, Lutzenberger W, Preissl H. Nouns and verbs in the intact brain: Evidence from event-related potentials and high-frequency cortical responses. *Cerebral Cortex*. 1999; 9(5):497–506. [PubMed: 10450894]
- Rauschecker JP, Scott SK. Maps and streams in the auditory cortex: Nonhuman primates illuminate human speech processing. *Nature Neuroscience*. 2009; 12(6):718–24.10.1038/nn.2331
- Sahin NT, Pinker S, Cash SS, Schomer D, Halgren E. Sequential processing of lexical, grammatical, and phonological information within broca's area. *Science*. 2009; 326(5951):445–9.10.1126/science.1174481 [PubMed: 19833971]
- Schoonover, CE. *Portraits of the mind : Visualizing the brain from antiquity to the 21st century* (illustrated ed.). New York: Abrams; 2010.
- Seung, S. *Connectome how the brain's wiring makes us who we are*. Boston: Houghton Mifflin Harcourt; 2012. Retrieved from <http://site.ebrary.com/id/10531983>
- Shukla M, White KS, Aslin RN. Prosody guides the rapid mapping of auditory word forms onto visual objects in 6-mo-old infants. *Proceedings of the National Academy of Sciences*. 2011; 108(15): 6038.
- Smolensky, P.; Legendre, Géraldine. *The harmonic mind : From neural computation to optimality-theoretic grammar*. Cambridge, Mass: MIT Press; 2006.
- Sporns, O. *Networks of the brain* (illustrated ed.). Cambridge, Mass: MIT Press; 2010.
- Telkemeyer S, Rossi S, Koch SP, Nierhaus T, Steinbrink J, Poeppel D, Wartenburger I, et al. Sensitivity of newborn auditory cortex to the temporal structure of sounds. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*. 2009; 29(47):14726–33.10.1523/JNEUROSCI.1246-09.2009 [PubMed: 19940167]
- Weinberg, S. *Lake views : This world and the universe*. Cambridge, Mass: Harvard University Press; 2009.
- Wernicke, C. *Der aphasische Symptomencomplex: Eine psychologische Studie auf anatomischer Basis*. Cohn & Weigert; 1874.
- Wilson, EO. *Consilience : The unity of knowledge* (reprint ed.). New York: Random House Digital; 1999.
- Yizhar O, Fenno LE, Davidson TJ, Mogri M, Deisseroth K. Optogenetics in neural systems. *Neuron*. 2011; 71(1):9–34. [PubMed: 21745635]
- Zatorre RJ, Gandour JT. Neural specializations for speech and pitch: Moving beyond the dichotomies. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2008; 363(1493):1087–1104.
- Zatorre RJ, Belin P, Penhune VB. Structure and function of auditory cortex: Music and speech. *Trends in Cognitive Sciences*. 2002; 6(1):37–46. [PubMed: 11849614]

Figure 1.
The organization of Broca's region, Amunts et al. (2010)

Figure 2.
Hickok & Poeppel 2007, dorsal and ventral streams

Experiment 1 (1.5T)

A. Topography (12 subjects)

Experiment 2 (replication at 3T)

C. Topography (8 males)

Figure 3.
Giraud et al. 2007, Neuron

Figure 4.
Perani et al. 2011, PNAS, dorsal stream development

Linguistics**Neuroscience*****Fundamental elements of representation (at a given analytic level)***

distinctive feature
syllable
morpheme
noun phrase
clause

dendrites, spines
neuron
cell-assembly/ensemble
population
cortical column

Fundamental operations on primitives (at a given analytic level)

concatenation
linearization
phrase-structure generation
semantic composition

long-term potentiation (LTP)
receptive field
oscillation
synchronization

Figure 5.
From Poeppel & Embick 2005