Linking language and cognition to neuroscience via computation

Report from the NSF Workshop held May 23-24 in Arlington VA

Submitted by David Poeppel, NYU
June 2013

Report Outline

1. Introduction

Understanding the human brain is of central importance to the biological, cognitive, and social sciences. Its structure and function lie at the basis of everything we do and experience. All basic bodily functions (heart, lung, temperature control, etc.) are controlled and regulated by the central nervous system, of course. But - more relevant in the present context - to *perceive*, *remember*, *feel*, *evaluate*, *decide*, *act*, *speak* … all are critical aspects of the mind and are functions that the brain instantiates. A theoretically motivated and nuanced, biologically sophisticated and detailed, and computationally explicit understanding of **how the brain implements the functions that constitute the mind will be of fundamental significance, for basic research, for translational applications, and for technological development**.

The small and highly focused workshop was constructed on the assumption that **the study of language and its neural foundations can and should play a pivotal role in making progress in the investigation of complex brain function**. Language research builds on a rich theoretical basis, established after several decades of research on the computations and representations that comprise language. Moreover, the neurobiological tools to study language are of increasingly high resolution and analytic sophistication. However, the computational basis of how the brain operates with linguistic representations remains underspecified and poorly understood. We are missing relevant computational analyses (at the right level of abstraction) to link language processing and neuroscience. As a consequence, the particular emphasis of the workshop was on computational linking hypotheses. The goal was and continued to be to identify **new directions in the 'computational neurobiology of language**.'

After briefly describing the structure of the workshop in Section 2, we describe the nature of the key intellectual challenge in Section 3 and make some general recommendations that arose in Section 4. In Section 5, a few comments are offered on the recent White House BRAIN initiative. Section 6 provides appendices.

2. Organization of the workshop

The workshop was convened by David Poeppel from NYU. The meeting was held at the NSF in Arlington VA on May 23-24, 2013. Keith Doelling from NYU, Dr. Poeppel's research assistant, helped with various aspects of the organization and acted as note-taker during the workshop.

Sixteen participants from a range of different disciplines attended (see Appendix 6.1). The participants were selected on the basis of their disciplinary expertise and their interdisciplinary savvy. The represented areas of expertise included **linguistics** (Berent, Berwick, Embick, Hornstein, Idsardi, Rapp, Smolensky, Sprouse), **computation** (Berwick, Embick, Gallistel, Heeger, Idsardi, Ma, Smolensky), **systems neuroscience** (Heeger, Hoffman, Ma, Melloni), and **cognitive neuroscience** (Berent, Embick, Heeger, Hickok, Idsardi, Melloni, Newport, Poeppel, Rapp).

The workshop was organized around a few longer presentations that summarized particular perspectives (*the linking problem*: Embick, Smolensky; *computation in vision*: Heeger; *computation in speech production*: Hickok) and sections comprised of shorter presentations on specific topics (see Appendix 6.2 for the full schedule). A considerable amount of time was dedicated to discussion.

Prior to the workshop, all participants submitted readings, either of a general nature or specific to their presentations. A reading list/syllabus was constructed, and these papers formed the basis both for the framing of the larger problem and for specific issues. The papers are listed in Appendix 6.3.

3. The nature of the problem

The experimental techniques available to characterize human brain function are becoming increasingly sophisticated and achieving ever-higher resolutions in space and time. Approaches ranging from hemodynamic (fMRI, PET, NIRS) and electrophysiological recordings (EEG, MEG, ECog) to stimulation and interference techniques (TMS, tDCS) – as well as the impressive new analytic tricks (multi-voxel pattern classification, ICA, oscillatory coherence analyses, etc.) - are being applied to virtually every aspect of human experience.

In the case of language research, all these approaches have been used to good effect. We now have quite reasonable 'brain maps,' at the level of gyral and sulcal anatomy, that specify where in the brain the major operations occur that underlie various aspects of language processing (e.g. Hickok & Poeppel 2007; Kotz & Schwartze 2010; Friederici 2012). Typical current models point to regions (e.g. Broca's region), streams (e.g. dorsal versus ventral), hemispheres, or networks of areas that are implicated in phonological processing, lexical access, syntactic analysis, and so on. Hickok in his presentation (Day 1) presented a well-developed functional anatomic model of speech production of this type. Indeed, in her comments (Day 2), Elissa Newport emphasized that analyses of brain activation data from hearing and signing speakers/listeners as well as data on brain plasticity in the context of abnormal development (e.g. hemispherectomy) point to a subtle balance between specificity for language versus plasticity, suggesting that some regions are specialized for different language functions.

Notwithstanding such successes in generating *spatial maps*, the relation between language and the brain is at best understood at a correlational level. That is to say, we do have good reasons to believe that, say, some *area x* relates systematically to some *language function y*, in some way. However, explanatory understanding is strikingly absent concerning how neural circuits might account for the implementation of the specific operations that underpin some linguistic computations and representations. This key point was made in detail in the opening presentation, on the linking problem, by David Embick (Day 1), an argument that we summarize in some detail below, as it provides a useful framework to appreciate many issues central to the meeting.

The mapping problem

Embick's point of departure was the 'mapping problem.' One way to make clear the considerable challenges that more integrated (and integrative) interdisciplinary research on brain and language faces is to be explicit about the conceptual infrastructure of the two domains that are to be linked. Addressing this linking or *mapping problem* - what is the relationship between the 'parts list' of language and the 'parts list' of neuroscience – is rather more difficult than it might seem at first glance, ultimately requiring the development of appropriate linking hypotheses between the different domains of study. Poeppel & Embick (2005) and Poeppel (2012) discuss this challenge in some detail, describing the 'contact' or 'interface problems' between linguistics and neuroscience and the diagnoses. Figure 1 below (from Embick's talk) illustrates the core of the mapping problem.

Language research develops computational-representational theories (CR), in which the primitive objects include 'distinctive feature' or 'morpheme' and the primitive operations include 'feature spreading', 'merge,' etc. Similarly, the neurosciences generate neurobiological theories and models (NB) in which the primitive objects include 'dendrite' or 'cortical column' and the primitive operations 'long-term potentiation' or 'oscillation.'

	CR	NB
Objects		
	distinctive feature	dendrite/spine
	timing slot	neuron
	morpheme	ensemble
	phrase	cortical column
Operations		
	feature spreading	long term potentiation (LTP)
	merge	oscillation
	concatenation	receptive field
	semantic composition	synchronization

(NB The expression *mapping* as used here does not refer to the spatial assignment of putative linguistic or psychological functions to brain areas, be they distributed or localized, microscopic or macroscopic. It is an empirical question whether such a spatial assignment of function works. Rather, mapping refers to the investigation of the necessary *formal relations* between two sets of hypothesized inventories, the inventory constructed by the language sciences and that constructed by the neurosciences.)

So how do the putative primitive units of analysis of CR models from language research (left column) map on to the putative primitive units of NB theories (right column)? The cognitive sciences, including linguistics and psychology, provide detailed analyses of the ontological structure of the domain (call this the 'cognome,' i.e. the comprehensive list of elementary mental representations and operations). Neurobiology provides a growing list of the available neural structures. To exemplify, 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 the computations and representations that constitute a speaker's knowledge of language, as well as 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? Intuitively simple reductionist attempts can be stated (say, e.g., 'neuron = syllable') but cannot be interpreted in any sensible way. Such 'naïve alignments' may seem pleasing at first glance - but that pleasure dissipates if one has to actually do the work of representation and computation to capture generalizations of interest to the language sciences. The fact of the matter is that we have very little (or really no) idea how the stuff of thought relates to the stuff of brains, in the case of speech and language – and in any other case (Gallistel & King 2011; Mausfeld 2012). There is no simple reduction/alignment of basic objects in the two domains; i.e. no linking theory between CR and NB domains (the "ontological incommensurability problem" of Poeppel & Embick 2005). So what then are the prospects for connections between CR and NB?

Linking hypotheses and computation

Embick's presentation suggested "the computational imperative" and developed three types of possible connections between CR and NB approaches. Embick argued, building on Marr and others, that linking language and neuroscience in an explanatory fashion will require the formulation of computationally explicit linking hypotheses - at the right level of abstraction. Figure 2 (below), from Carandini (2012), depicts the same logical problem as discussed by Embick, described by a computational neuroscientist focusing on vision. The approach advocated by Carandini builds on the ideas of Marr (1982), the approach that is also championed for language studies by Poeppel et al. (2008). To connect neural circuits (the objects of neuroscience) to behavior (the objects of the cognitive sciences) requires computational hypotheses.

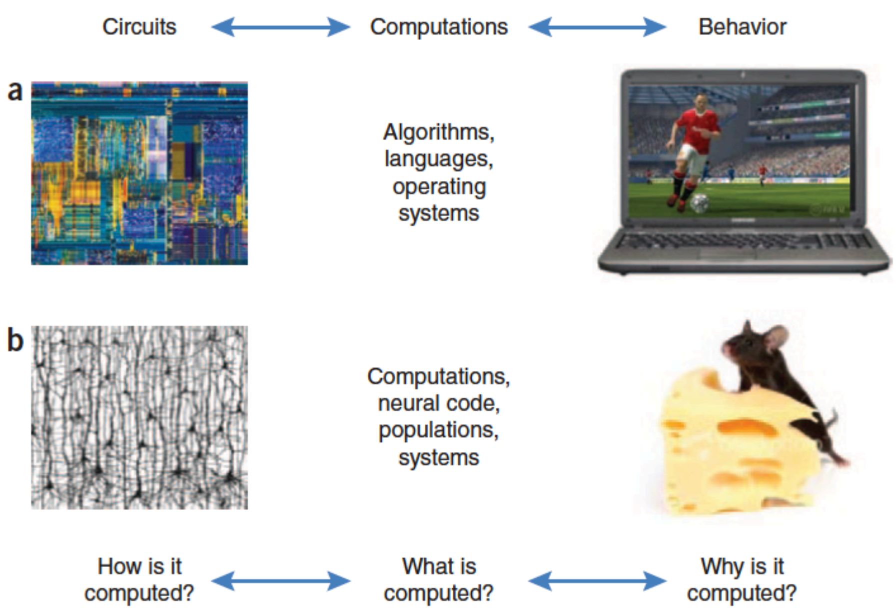

(**a**) The wiring of a fraction of an Intel microprocessor and a laptop playing a popular videogame (FIFA 12). (**b**) Pyramidal neurons in cortex (detail of a drawing by Ramon y Cajal) and a mouse engaged in a pleasant behavior.

Embicks "computational imperative" builds on the following idea: CR theories of computational and algorithmic levels of language (i.e. "theoretical linguistics", "psycholinguistics") must be used to investigate how the brain encodes/represents/computes language. Herein he follows the general assumption of the centrality of computational theory for exploring the brain, a topic expounded at length in Gallistel and King (2009).

More specifically, Embick suggests that CR (linguistic) analyses decompose phenomena into subroutines (fundamental operations/representations); these articulated CR theories should inform the investigation of NB structures (whether in conjunction with neuroimaging or neuropsychological methods). Critically, using CR theory to investigate how the brain computes language is not the same as using NB to "validate" CR theory; this is impractical given the different levels of detail currently achieved in each domain (the "granularity mismatch problem" of Poeppel & Embick 2005). CR theories of language are highly articulated, NB theories of language are, in most domains, less articulated.

Three types of CR-NB connections are raised in this context. **(1) CR Neurolinguistics**: CR theories of language are used to investigate the what and how of the NB side of language. CR neurolinguistics would be a huge achievement. In order to get a working theory of how the brain computes, we would have to answer questions about encoding etc. that are at the frontiers of research (at least for the more abstract parts of language). **(2) Integrated Neurolinguistics**: this would amount to CR neurolinguistics plus, critically, the NB data provide crucial evidence that adjudicates among different CR theories. **(3) Strong Neurolinguistics**: (1+2) AND some feature of an NB structure/function explains WHY the CR theory of language involves particular computations and representations (and not others). The idea behind ttrong neurolinguistics is as follows: in order to truly explain some aspects of language, it is necessary to look at NB, because NB explains why we find what we find in CR. A key notion here is what it might mean for NB to be specialized for computations or representations.

Naturally, the goal is to achieve the third type of theory, that is to say an account of a complex cognitive phenomenon in which its subtle features are explained by a detailed neurobiological accounts of the structure. The aspirational goal has to be to develop a mechanistic understanding of neural computation, but this is immensely challenging in any domain of perception and cognition (Denk et al. 2012, *Nat Rev Neurosci*). Given the state-of-the-art, progress of the first or second type must already be considered substantial. It is worth noting that an attempt to provide a soup-to-nuts, fully linked account was provided by Smolensky (Day 2 lecture). The extent to which the proposal was plausibly grounded in NB type theories was debated; however, the approach demonstrates how an explicit set of links can be drawn from a linguistic theoretical question to well motivated computations and ultimately to the level of circuits.

4. Recommendations

With these issues in mind, the two-day workshop explored the idea that computational analyses at a certain level of abstraction and granularity might form appropriate linking hypotheses between the 'alphabets' of linguistic research and neurobiology, adopting the playbook of Marr and others. If this is on the right track, it opens up new experimental and theoretical directions in cognitive neuroscience in general, and brain and language research, in particular. Such a computational perspective also has clear implications for how we discuss the evolution and acquisition of language, moving from high-level descriptions ('evolving syntax') to computationally explicit operations (say 'developing a circuit that executes linguistic operations like those that combine or concatenate constituents') that may or may not be specialized and are attributes of neural circuits.

Of the wide range of suggestions raised in the talks and discussions, we select here four suggested ideas at different levels of grain size, stimulated by considerations from different disciplines. All deal with **the challenge of finding new and tractable linking hypotheses**.

• *Computer science*: Capitalize on what has been learned in computer architecture in seeking links to neural computation.
• *Linguistics*: Use recent discoveries from linguistics to study domain general and language specific basic operations.
• *Cognitive neuroscience*: Functional anatomic models of speech perception and production implicate operations like "coordinate transformation." Such concepts can benefit from extensive physiology and modeling in visual neuroscience.
• *Genetics*: Embrace contemporary genetic approaches and identify links to language processing and individual differences.

4.1. As discussed by Embick as well as several others (Berwick, Idsardi, Gallistel), the computational perspective developed by Marr has played an important role in framing the questions. Berwick in his talk (Day 1) developed the argument how to build on these ideas: ***leveraging the modern computer science approach to complex computational architectures***.

There seem to be at least two "grand challenge areas" that call for a renewed investigation of how our knowledge of computation might inform the neurobiology-language connection, the first having to do with our understanding of how complex, parallelized computations work, based our increased understanding of computer architecture; and the second concerning the way in which memory might be organized biologically, via content-addressable systems as opposed to conventional random access memory. Each issue is considered in turn.

(i) First, over the past several decades, computer science has developed an extremely powerful set of methods for designing, analyzing, and testing very complex computational devices that work in a highly parallel way. Modern CPU design takes these abstract tools for granted, including dataflow analysis, models for pipelined computer architecture, and various 'layers' of abstract machines, all organized with the understanding that many operations take place concurrently, e.g., fetching several instructions to execute while decoding and executing many others; routing operands and operators in parallel; and so forth, all done so as to take into account complex flow and race conditions. The design and software tools in this area constitute a mature technology, as exemplified in now-standard texts such as Hennessey and Patterson's *Computer Architecture, a Quantitative Approach*. As the left picture of the data paths in a modern CPU illustrate, the resulting devices have begun to resemble the complex wirings seen in neuroscience, both in structure and function.

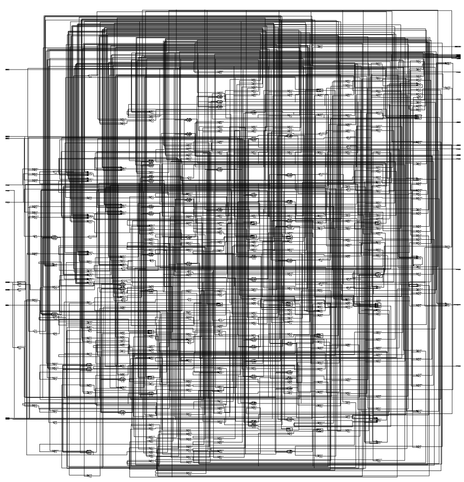

Example of modern highly-pipelined computer central processing unit (CPU).

This immediately suggests the following forward-looking grand challenge:

• Determine whether we can leverage our understanding of computer architecture in the context of the neurobiology of language.
• In particular, can we develop a 'parts list' and a set of abstract machine layers that go beyond the sorts of circuits posited in neuroscience (oscillators, etc.), and beyond stack and register devices as suggested in Gallistel and King, to begin to address the kinds of computations that take place in neurobiology?

(ii) Second, it has often been suggested that brains might use a memory system that is quite different from conventional computers that arrange information according to numbered addresses in memory (a random access memory), instead using a system where memory is organized by its content. This sort of memory has, in fact, been analyzed in computer science but has rarely been implemented, since random access memory has generally been far less expensive. In this case, it becomes extremely fast to pick out all the elements that share a particular property. Given the direction of some current linguistic theories that are grounded on using the properties of word features as labels, this kind of representation becomes more attractive. For instance, one of the basic operations in such a system consists in taking two words, like *red* and *apple*, and then using one of them, say, *apple* to label the unit *red apple*, with *apple* serving as an associative 'key' to access the entire unit. If so, it would seem appropriate at this time to examine how such **content-addressable systems** might be developed to dovetail with these new linguistic accounts along with psycholinguistic results that have at times seemed at odds with explicit representation of "tree structure." Here too, there is a base of computational work to draw on, e.g., as in Oldfield, Williams, and Wiseman (1987; Content-addressable memories for storing and processing recursively subdivided trees. Electronic Letters, 1987, 23:6.) A first goal would be to implement current linguistic theories using these sorts of representational methods rather than 'conventional' random-access approaches.

4.2 ***Syntactic primitives***. The goals of syntactic research over the last 20 years have re-aligned nicely with the goals of cognitive and systems neuroscience: to identify fundamental neuronal computations that (i) underlie a large number of (linguistic) phenomena, and (ii) rely as little as possible on domain-specific properties. As a concrete example, the syntactic theory known as Minimalism has formulated a two-step syntactic function called Merge that separates nicely into a domain-general neuronal computation that *combines* elements (somewhat akin to *binding*, in the context of systems neuroscience), and a potentially more domain-specific computation that *labels* the output of the binding computation:

(1) Bind: Given an expression A and an expression B, bind A,B → {A,B}
(2) Label: Given a combined {A,B}, label the complex A or B; → {A A,B} or {B A,B}

Much recent work in Minimalism has suggested that many of the complex properties of natural languages can be modeled as repeated applications of the Bind and Label computations. Furthermore, the formal characterization of these computations in set-theoretic terms provides a computational level description that is similar to the formal characterization of neuronal computations in other domains of cognition (e.g., normalization functions; Carandini & Heeger, 2011). This new direction in syntactic theory marks a radical departure from earlier syntactic theories, which contained a large number of disparate kinds of rules, and relied heavily on domain-specific properties of those rules. As such, the time seems right for a renewed collaboration between syntacticians and cognitive/systems neuroscientists in a search for the neural circuits that subserve fundamental syntactic computations like *Combine/Bind* and *Label*, and ultimately the neuronal encoding of the computations themselves. Cognitive and systems neuroscientist who were previously dissuaded by the myriad rules of earlier syntactic theories may be heartened to learn that syntacticians have already begun to reformulate syntactic theories in terms of fundamental neuronal computations.

4.3 ***Coordinate transformations.*** Large-scale functional anatomic models of speech perception and speech production are converging on problems that have been addressed more thoroughly in other parts of the neurosciences. This is a convergence that ought to be taken advantage of in terms of computational models and neurophysiological characterization. One specific point of contact concerns so called coordinate transformations. In the literature on motor control, visually guided reaching, eye movements, and so on, this concept is well developed and extensively studied. The underlying idea and logic has recently been imported into speech and language research (e.g. Hickok & Poeppel 2007; Hickok 2012; Tian & Poeppel 2013). The intuition is captured by the simple observation that acoustic input and motor output are calculated in different coordinate systems (i.e. acoustic space versus motor/joint space). The translation between such representations will illuminate questions about perception and action.

4.4 ***Language, genetics, and neuroscience***. Three recent innovations in scientific methodology suggest a 'grand challenge' to link language science to the human genome and to neuroscience: (i) so-called "next-generation" sequencing (NGS), which will lower the cost of whole genome sequencing so that individuals can be sequenced for quite modest costs; (ii) new, psychometrically-well grounded analysis of individual linguistic knowledge that can be applied to large subpopulations; and (iii) neuroscience techniques using genetic tagging to visualize active in vivo neural circuitry. The 'grand challenge' goal would be to arrive at a fuller characterization of the human language phenotype and its natural variation in human populations, in order to begin to establish better connections between human language and its genetics, as well as its neurobiological realization. We cover the first two of these three methodological advances in turn. The third approach is gaining traction in animal preparations but it is not yet clear how it may or may not apply to language studies in humans, and so we leave it for a future discussion.

1. *NGS and identifying recent selection in the human genome*. With the advent of NGS and the 1000 genomes project (http://www.1000genomes.org/), it has now become practicable to engage in genome-wide scans for areas that have undergone positive selection in the past 30-60 thousand years (and possibly longer); which is to say, those genomic properties that make us uniquely human. In a recent paper in Cell (Identifying recent adaptations in large-scale genomic data, 14 February 2013, 703-714), Grossman et al. demonstrate that one can retrieve upwards of 400+ regions that reveal such a signature of positive selection. They then proceed to demonstrate the power of their method by selecting one previously unknown candidate region (an inter-genic region), and establishing experimentally the corresponding phenotype, the reduction of septic shock due to bacterial infection. This is a sensitive, powerful technique of broad applicability to locate the actual adaptive traits that match with genetic changes in the human lineage.

2. *Uncovering the human language phenotype*. To be able to use genomic information, we need to know more about the human language phenotype and its variation in individuals and subpopulations. Though it may be possible to find Mendelian and near-Mendelian language-related traits, as with FOXP2 and oral dyspraxia, it seems more likely that language is a complex, multifactorial trait. The first step is finding out more about human language variation, in the biological sense. The challenge of finding human variation in language ability can now be aided by the development by Jon Sprouse (see presentation on Day 1) and others for the application of well-established psychometric tests to human sentence judgments, the mainstay of modern linguistic theory. Sprouse and colleagues have established that sentence judgments are as reliable as any other standard psychometric test. In addition, they have shown how to construct such tests using Mechanical Turk to obtain large sample size assessments of linguistic data. Given this ability, a natural next step is to begin to assess the variation within subpopulations with respect to linguistic judgments, under various conditions, so as to begin the characterization of human variation within, e.g., syntactic abilities that have been well-established by linguistic theory. Given proper (substantial) sample sizes, one can then correct for confounds of several kinds. If this can be coupled with genomic scans of the relevant subpopulations, we may well be able for the first time to uncover variation in the human language phenotype, presumably a complex trait subject to analysis by the method of quantitative trait loci (QTL) analysis. While this approach has been broached occasionally in the literature (e.g. by Mabel Rice, Ken Wexler, and others), it has not been possible to carry out on a wide scale due to the lack of proper testing methodology and genomic analysis, both of these being bottlenecks that may now have been overcome.

5. White House BRAIN Initiative

A section of the workshop was dedicated to discussing the new initiative on neuroscience that has been publicized as a recent White House initiative. We briefly summarize some reactions here (summary comments courtesy of workshop participant Wei Ji Ma).

What the workshop group considered promising in the BRAIN Initiative:

- The BRAIN Initiative puts neuroscience, a field that has been one of the most successful interfaces between science and society, into the spotlight.
- The BRAIN Initiative is a much more promising concept than the European Union's "Human Brain Project" of large-scale brain simulation.
- The BRAIN Initiative recognizes that the study of the brain is an inherently interdisciplinary effort.

Some critical reactions:

- **Overemphasis on technology.** New experimental methods can go only so far. Needed are targeted experiments, distilled to conceptual understanding by theory and modeling. Theory and modeling are more than data mining. The BRAIN Initiative as anticipated in the well publicized Neuron paper is like trying to assemble Tycho Brahe's star catalogs without looking for Kepler's or Newton's laws.
- **Behavior matters.** Understanding the brain is impossible without understanding what the organism does in the world: perception, movement, language, thought. The BRAIN Initiative must connect to the SBE sciences and their efforts to relate brain to behavior.
- **Unrealistic goals can backfire.** If BRAIN produces terabytes of data without bringing us closer to understanding the brain or curing disease (for example because theory is lacking), the public might blame the scientists for failing to live up to the promise, which can jeopardize future support for basic research.

Some suggestions:

- **N for Neuro-Theory.** The N in BRAIN currently stands for Neuro-*technology*. We advocate to let it stand for "Neurotechnology and Neurotheory", for the reasons outlined above.
- **A scaffolding theme.** We suggest to give the BRAIN initiative a more substantive goal by choosing a broad scaffolding theme that connects researchers in different disciplines and at different levels of analysis (including behavior). Such a theme could help us refocus on the scientific questions rather than the technology.
- **The changing brain.** One example of such a theme could be "learning and memory". These are of great interest to the general public and could be related to societal issues like lifelong education and cognitive fitness. A tagline could be "The changing brain", and a realistic goal could be to "discover laws of how the brain learns and remembers" (note: laws, not *the* laws).
- **Interdisciplinary training.** Progress in brain research depends on researchers – both collectively and individually – being trained in multiple disciplines, ranging from biology, psychology, and cognitive science, to computer science, mathematics, engineering, and physics. This can be achieved through additional IGERT funding targeting neuroscience specifically, as well as support for developing similar programs at the undergraduate level.

6. Appendices

Appendix 6.1 Participants

Guest	Affiliation	Website
Iris Berent	Northeastern	http://www.northeastern.edu/berentlab/
Bob Berwick	MIT	http://bcs.mit.edu/people/berwick.html
David Embick	U Penn	http://www.ling.upenn.edu/~embick/
Randy Gallistel	Rutgers	http://ruccs.rutgers.edu/~galliste/index.php/
David Heeger	NYU	http://www.cns.nyu.edu/heegerlab/
Greg Hickok	UC Irvine	http://alns.ss.uci.edu
Kari Hoffman	York University	http://www.yorku.ca/khoffman/
Norbert Hornstein	University of Maryland	http://ling.umd.edu/~hornstein/
Bill Idsardi	University of Maryland	http://ling.umd.edu/people/person/bill-idsardi/
Weiji Ma	Baylor	http://neuro.bcm.edu/malab/
Lucia Melloni	Max Planck Institute	http://brain.mpg.de/?id=503
Elissa Newport	Georgetown University	http://cbpr.georgetown.edu/faculty/elissa_newport/
David Poeppel	NYU	http://psych.nyu.edu/clash/poeppellab
Brenda Rapp	Johns Hopkins	http://cogsci.jhu.edu/labs/CogNeuro/
Paul Smolensky	Johns Hopkins	http://cogsci.jhu.edu/people/smolensky.html
Jon Sprouse	UC Irvine	http://www.socsci.uci.edu/~jsprouse/
Keith Doelling	NYU	http://psych.nyu.edu/clash/poeppellab/theLab?name=keith

Appendix 6.2 Program

Thursday, May 23

8:30 am Coffee+
8:45 am Welcome and opening remarks, David Poeppel
9:00 am Welcoming Remarks by Myron Gutmann (SBE AD)
9:15 am Our list of questions - DP
9:30 am Dave Embick - "The linking problem"
10:15 am Greg Hickok - "Speech production and computational foundations"
11:00 am Coffee Break
11:15 am David Heeger - "Motion analysis and canonical computations"

12:00 pm Lunch

Syntactic Computation

1:30 pm Norbert Hornstein
2:00 pm Jon Sprouse
2:30 pm Bob Berwick
3:00 pm Coffee

Systems Neuroscience, Computational Analysis

3:30 pm Kari Hoffman
4:00 pm Wei Ji Ma
4:30 pm Lucia Melloni

5:00 pm R & R
7:00 pm Dinner (at a private dining room in Pinzimini inside the Westin Arlington)

Friday, May 24

8:30 am Coffee+
8:45 am The White House BRAIN Initiative
9:30 am Paul Smolensky - "The linking problem, another perspective"
10:15 am Coffee
10:30 am Elissa Newport

Sounds and words, heard/spoken/read/written

11:00 am Bill Idsardi
11:30 am Iris Berent
12:00 pm Brenda Rapp

12:30 pm Lunch

2:00 pm A la carte discussion
3:00 pm Randy Gallistel
4:00 pm Closing Remarks and Adjourned

Appendix 6.3 Reading List/References

Preparatory Readings

Some of the readings that motivated this workshop

1. Poeppel, D. & Embick, D, 2005
2. Poeppel, D., 2012
3. Carandini, M., 2012
4. Marr, D., 1982
5. Mausfeld, R., 2012

Readings from the linguistic perspective

6. Hornstein, N. 2013
7. Berwick, R. Friederici, A., Chomsky, N., & Bolhuis, J., 2012.
8. Prince, A. & Smolensky, 1997
9. Rizzi, L., 2011
10. Poeppel, D., & Idsardi, B., 2011
11. Berent, I., in press.
12. Smolensky & Legendre, 2006 - Excerpt from *The Harmonic Mind*
13. Smolensky, 2006 - Section 23.1 from *The Harmonic Mind*

Readings from the neurobiological perspective

14. Carandini, M. & Heeger, D., 2011
15. Hickok, G., 2012
16. Friederici, A., 2011
17. Pouget, A., Deneve, S., & Duhamel, J., 2002
18. Turesson, H.K., Logothetis, N.K., & Hoffman, K.L., 2012
19. Melloni, L., Schwiedrzik, C.M., Müller, N., Rodriguez, E., & Singer, W., 2011
20. Ma, W.J., Navalpakkam, V., Beck, J., van den Berg, R., & Pouget, A., 2011
21. Kriegeskorte, N., 2009

Other exciting readings to consider

22. Giurfa, M., Zhang, S., Jenett, A., Menzel, R. & Srinivasan, M., 2001
23. Weber, J., Peterson, B.K., & Hoekstra, H.E., 2013

Papers by Workshop Section:

Embick - "The linking problem":

- Poeppel, D. & Embick, D, 2005
- Poeppel, D., 2012
- Carandini, M., 2012
- Marr, D., 1982

- Mausfeld, R., 2012

Hickok - "Speech production and computational foundations"

- Hickok, G., 2012

Heeger - "Motion analysis and canonical computations"

- Carandini, M. & Heeger, D., 2011
- Pouget, A., Deneve, S., & Duhamel, J., 2002

Syntactic Computation

- Hornstein, N. 2013
- Berwick, R. Friederici, A., Chomsky, N., & Bolhuis, J., 2012.
- Rizzi, L., 2011
- Friederici, A., 2011

Systems Neuroscience, Computational Analysis

- Ma, W.J., Navalpakkam, V., Beck, J., van den Berg, R., & Pouget, A., 2011
- Turesson, H.K., Logothetis, N.K., & Hoffman, K.L., 2012
- Melloni, L., Schwiedrzik, C.M., Müller, N., Rodriguez, E., & Singer, W., 2011

Paul Smolensky - "The linking problem, another perspective"

- Prince, A. & Smolensky, 1997
- Smolensky & Legendre, 2006 - Excerpt from *The Harmonic Mind*
- Smolensky, 2006 - Section 23.1 from *The Harmonic Mind*

Sounds and words, heard/spoken/read/written

- Berent, I., in press.
- Giurfa, M., Zhang, S., Jenett, A., Menzel, R. & Srinivasan, M., 2001
- Poeppel, D., & Idsardi, B., 2011
- Kriegeskorte, N., 2009

Randy Gallistel

- Weber, J., Peterson, B.K., & Hoekstra, H.E., 2013

Berent, I. (2013). The phonological mind. *Trends Cogn Sci*. doi: 10.1016/j.tics.2013.05.004

Berwick, Robert C., Friederici, Angela D., Chomsky, Noam, & Bolhuis, Johan J. (2013). Evolution, brain, and the nature of language. *Trends in cognitive sciences, 17*(2), 89-98.

Carandini, M. (2012). From circuits to behavior: a bridge too far? *Nat Neurosci, 15*(4), 507-509. doi: 10.1038/nn.3043

Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nat Rev Neurosci, 13*(1), 51-62. doi: 10.1038/nrn3136

Friederici, A. D. (2011). The brain basis of language processing: from structure to function. *Physiol Rev, 91*(4), 1357-1392. doi: 10.1152/physrev.00006.2011

Giurfa, M., Zhang, S., Jenett, A., Menzel, R., & Srinivasan, M. V. (2001). The concepts of 'sameness' and 'difference' in an insect. *Nature, 410*(6831), 930-933. doi: 10.1038/35073582

Hickok, G. (2012). Computational neuroanatomy of speech production. *Nat Rev Neurosci, 13*(2), 135-145. doi: 10.1038/nrn3158

Hornstein, N. (Forthcoming). Three grades of grammatical involvement: syntax from a minimalist perspective. . *Mind and Language*.

Kriegeskorte, N. (2009). Relating Population-Code Representations between Man, Monkey, and Computational Models. *Front Neurosci, 3*(3), 363-373. doi: 10.3389/neuro.01.035.2009

Ma, W. J., Navalpakkam, V., Beck, J. M., Berg, Rv, & Pouget, A. (2011). Behavior and neural basis of near-optimal visual search. *Nat Neurosci, 14*(6), 783-790. doi: 10.1038/nn.2814

Marr, David. (1982). *Vision*: W.H.Freeman & Co Ltd.

Mausfeld, R. (2012). On some unwarranted tacit assumptions in cognitive neuroscience. *Front Psychol, 3*, 67. doi: 10.3389/fpsyg.2012.00067

Melloni, L., Schwiedrzik, C. M., Muller, N., Rodriguez, E., & Singer, W. (2011). Expectations change the signatures and timing of electrophysiological correlates of perceptual awareness. *J Neurosci, 31*(4), 1386-1396. doi: 10.1523/JNEUROSCI.4570-10.2011

Poeppel, D. (2012). The maps problem and the mapping problem: two challenges for a cognitive neuroscience of speech and language. *Cogn Neuropsychol, 29*(1-2), 34-55. doi: 10.1080/02643294.2012.710600

Poeppel, D., & Embick, D. (2005). Defining the Relation Between Linguistics and Neuroscience. In A. Cutler (Ed.), *Twenty-first century psycholinguistics: Four cornerstones* (pp. 103-118). Mahwah, NJ: Lawrence Erlbaum Assoc.

Poeppel, David, & Idsardi, William. (2011). Recognizing words from speech: The perception-action-memory loop. In G. Gaskel & P. Zwitserlood (Eds.), *Lexical Representation: a multidisciplinary approach* (pp. 171-196). Berlin: de Gruyter Mouton.

Pouget, A., Deneve, S., & Duhamel, J. R. (2002). A computational perspective on the neural basis of multisensory spatial representations. *Nat Rev Neurosci, 3*(9), 741-747. doi: 10.1038/nrn914

Prince, A., & Smolensky, P. (1997). Optimality: from neural networks to universal grammar. *Science, 275*(5306), 1604-1610.

Rizzi, Luigi. (2012). Core linguistic computations: How are they expressed in the mind/brain? *Journal of Neurolinguistics, 25*(5), 489-499.

Smolensky, P., & Legendre, G. (2006). *The harmonic mind: from neural computation to optimality-theoretic grammar.* . Cambridge, MA: MIT Press.

Turesson, H. K., Logothetis, N. K., & Hoffman, K. L. (2012). Category-selective phase coding in the superior temporal sulcus. *Proc Natl Acad Sci U S A, 109*(47), 19438-19443. doi: 10.1073/pnas.1217012109

Weber, J. N., Peterson, B. K., & Hoekstra, H. E. (2013). Discrete genetic modules are responsible for complex burrow evolution in Peromyscus mice. *Nature, 493*(7432), 402-405. doi: 10.1038/nature11816